Prepared for submission to JHEP

arXiv:1403.5215v1 [math.GT] 20 Mar 2014

A Spectral Perspective on Neumann-Zagier

Tudor Dimofte[1] **Roland van der Veen**[2]

[1] *Institute for Advanced Study, Einstein Dr., Princeton, NJ 08540, USA*

[2] *Korteweg-de Vries Institute for Mathematics, University of Amsterdam P.O. Box 94248 1090 GE Amsterdam The Netherlands*

Abstract: We provide a new topological interpretation of the symplectic properties of gluing equations for triangulations of hyperbolic 3-manifolds, first discovered by Neumann and Zagier. We also extend the symplectic properties to more general gluings of $PGL(2,\mathbb{C})$ flat connections on the boundaries of 3-manifolds with topological ideal triangulations, proving that gluing is a K_2 symplectic reduction of $PGL(2,\mathbb{C})$ moduli spaces. Recently, such symplectic properties have been central in constructing quantum $PGL(2,\mathbb{C})$ invariants of 3-manifolds. Our methods adapt the spectral network construction of Gaiotto-Moore-Neitzke to relate framed flat $PGL(2,\mathbb{C})$ connections on the boundary $\mathcal{C}$ of a 3-manifold to flat $GL(1,\mathbb{C})$ connections on a double branched cover $\Sigma \to \mathcal{C}$ of the boundary. Then moduli spaces of both $PGL(2,\mathbb{C})$ connections on $\mathcal{C}$ and $GL(1,\mathbb{C})$ connections on Σ gain coordinates labelled by the first homology of Σ, and inherit symplectic properties from the intersection form on homology.

Contents

1 Introduction and motivation **2**
- 1.1 Symplectic structures from homology 4
- 1.2 Gluing 6
- 1.3 Example: knot complement 8

2 Preliminaries **11**
- 2.1 Odd homology 11
- 2.2 Twisted homology 12
- 2.3 Boundaries of framed 3-manifolds 13
- 2.4 The canonical cover 15
- 2.5 Gluing of canonical covers 17

3 Paths on the small boundary **20**
- 3.1 Paths and generators for $H_1^-(\Sigma)$ 20
- 3.2 Generators of $H_1^-(\Sigma)$ and intersection form 23
- 3.3 $\widetilde{G}$, $\widetilde{K}$, and the gluing map 26
- 3.4 Detour: Neumann's chain complex 28

4 Framed flat connections on boundaries **30**
- 4.1 Moduli spaces 31
- 4.2 Coordinate functions 32
- 4.3 Poisson brackets, symplectic form, and K_2 form 36
- 4.4 Gluing $PGL(2)$ connections on boundaries 38
- 4.5 Symplectic properties of the gluing equations 40
- 4.6 Example: Thurston's gluing equations, K_2 forms, and volumes 42

5 Non-abelianization **45**
- 5.1 Abelian flat connections 46
- 5.2 Gluing abelian connections 48
- 5.3 Spectral networks and non-abelianization 49
- 5.4 Non-abelianization commutes with gluing 52

A Odd results **53**
- A.1 Basics 53
- A.2 Chain complexes 55
- A.3 Cellular description 56
- A.4 Cutting and gluing 58

B Reconstructing framed flat connections **60**

1 Introduction and motivation

Systems with non-abelian gauge symmetry can sometimes be analyzed very effectively using related systems with abelian gauge symmetry. A famous example involves the use of (abelian) Seiberg-Witten theory in four dimensions to compute (non-abelian) Donaldson invariants [1]. Donaldson and Seiberg-Witten theory are smoothly connected in physics: they are two limits of the same four-dimensional quantum field theory.

In the present paper we use the same basic philosophy to study moduli spaces of flat (non-abelian) $PGL(2,\mathbb{C})$ connections in two and three dimensions, by means of closely related — but conceptually much simpler — moduli spaces of (abelian) $GL(1,\mathbb{C})$ flat connections. Again, these two types of moduli spaces occur naturally in the same physical systems. Moduli spaces of flat connections on a two-dimensional surface describe vacua of four-dimensional $\mathcal{N}=2$ supersymmetric theories of "class $\mathcal{S}$" [2, 3] (further developed in, *e.g.* [4–6]); while moduli spaces of flat connections on a 3-manifold describe vacua of three-dimensional $\mathcal{N}=2$ theories of "class $\mathcal{R}$" [7–10]. Just as 3-manifolds can have 2-dimensional boundaries, the 3d theories of class $\mathcal{R}$ (labelled by 3-manifolds) describe boundary conditions for the 4d theories of class $\mathcal{S}$ (labelled by surfaces).

The main question that we address, and hope to shed light on, is the symplectic nature of Thurston's gluing equations [11] for ideal triangulations of hyperbolic manifolds, and their generalizations. The gluing equations for cusped hyperbolic manifolds were first shown to have symplectic properties by Neumann and Zagier [12]. The symplectic properties immediately implied a formula for the variation of the volume of a hyperbolic 3-manifold as cusps are deformed. They have since been used to show that A-polynomials of hyperbolic 3-manifolds M are K_2-Lagrangian submanifolds in natural K_2-symplectic spaces associated to ∂M [13, 14], and that the hyperbolic structures on 3-manifolds can be systematically quantized [15–20].[1] They also played a crucial role in the construction of 3d $\mathcal{N}=2$ quantum field theories associated to 3-manifolds [10, 27], which (in principle) provide a categorification of hyperbolic invariants along the lines of [28].

The gluing equations and their symplectic properties have been generalized in many ways since the work of Neumann and Zagier. Neumann [29] showed that they held for topological ideal triangulations (not necessarily of hyperbolic manifolds). It later became clear that in the topological setting the gluing equations naturally describe a gluing of framed flat $PGL(2,\mathbb{C})$ connections (*e.g.*, [30, 31]), which include hyperbolic metrics. Symplectic properties of gluing equations were conjectured in [32] for 3-manifolds partially glued from ideal tetrahedra, with a

[1] Ideas about quantization of A-polynomials go back to [21, 22]. Alternative methods of quantization include skein calculus (*e.g.* [23]) and topological recursion (*e.g.* [24–26]), which are expected to be equivalent to triangulation constructions.

proposed proof in [33].[2] Symplectic properties were also conjectured for generalized $PGL(K)$ gluing equations in [31, 35], with recent proposed proofs in [33, 36].

Unfortunately, so far, all proofs of symplectic properties of gluing equations have involved subtle combinatorics, and have been relatively unintuitive. (The impressive works of [29, 36] are testament.) We seek to remedy this situation with an elementary topological construction.

First, we observe that symplectic properties of gluing equations have to do with framed flat connections on *boundaries* rather than interiors of 3-manifolds — for example, boundaries of ideal tetrahedra and the torus boundary of a fully-glued cusped 3-manifold. Thus, for the most part, we are dealing with an intrinsically two-dimensional problem. Then we borrow (and extend) a construction of Gaiotto, Moore, and Neitzke [6, 37] who showed, in the context of 4d $\mathcal{N} = 2$ theories of class $\mathcal{S}$, that the moduli space of framed flat $GL(K, \mathbb{C})$ connections with certain singularities on a surface $\mathcal{C}$ is (roughly) symplectomorphic to a space of flat $GL(1, \mathbb{C})$ connections on a K-fold branched cover $\Sigma \xrightarrow{\pi} \mathcal{C}$.[3] This correspondence was called a non-abelianization map. In [6], the cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ is a spectral cover, and the non-abelianization map was defined using the data of a related "spectral network" on $\mathcal{C}$ — hence the title of our paper. For our purposes, we will treat spectral networks (and the non-abelianization maps they induce) as purely topological objects.

The space of $GL(1, \mathbb{C})$ connections on Σ is extremely simple. It has coordinates $x_\gamma \in \mathbb{C}^*$ labelled by cycles $\gamma \in H_1(\Sigma, \mathbb{Z})$, and a Poisson bracket given by the intersection form in homology, $\{x_\gamma, x_{\gamma'}\} = \langle \gamma, \gamma' \rangle x_\gamma x_{\gamma'}$. Physically, $H_1(\Sigma)$ is just the electric-magnetic charge lattice of the 4d $\mathcal{N} = 2$ gauge theory labelled by a "UV curve" $\mathcal{C}$ and a Seiberg-Witten curve Σ. Via non-abelianization, the space of framed flat $GL(K, \mathbb{C})$ connections on $\mathcal{C}$ inherits the coordinates x_γ and their simple Poisson bracket, which coincides with the inverse of the Atiyah-Bott symplectic form [39]. The non-abelianization map can further be modified to provide a symplectomorphism between $PGL(K, \mathbb{C})$ connections on $\mathcal{C}$ and a projectivized space of $GL(1, \mathbb{C})$ connections on Σ, whose coordinates are labelled by elements of odd homology $\gamma \in H_1^-(\Sigma, \mathbb{Z}) := \ker \left[\pi_* : H_1(\Sigma) \to H_1(\mathcal{C})\right]$.

In special cases, the x_γ coincide with Fock-Goncharov cluster coordinates on spaces of framed flat connections, which complexify Thurston's shear coordinates (or, dually, Penner's length coordinates) in Teichmüller theory. As pointed out in [37, 40] and as we discover here, they can also be much more general. Even in the $PGL(2, \mathbb{C})$ case, the coordinates x_γ include complexified Fenchel-Nielsen coordinates on boundaries of 3-manifolds, of the type discussed in [9, 32, 35, 41, 42].

We claim that symplectic (and in fact K_2) properties of gluing equation are an obvious consequence of the topological fact that if a 3-manifold M is glued out of ideal tetrahedra $\{\Delta_i\}_{i=1}^N$ (or, in fact, other 3-manifolds) then the odd homologies of appropriate double-covers

[2] This situation is closely related to Bonahon's formula for the deformation of the volume of hyperbolic 3-manifolds with geodesic boundary [34].

[3] Such a spectral-cover construction of local systems also features (independently) in yet-unpublished work of Goncharov and Kontsevich [38], where coordinates on the relevant $GL(1, \mathbb{C})$ and $GL(K, \mathbb{C})$ moduli spaces are promoted to fully non-commutative variables.

of boundaries $\Sigma \xrightarrow{\pi} \mathcal{C} = \partial M$ and $\Sigma_{\Delta_i} \xrightarrow{\pi} \mathcal{C}_\Delta = \partial\Delta$ are related by a lattice symplectic reduction,

$$H_1^-(\Sigma, \mathbb{Z}) \simeq \oplus_{i=1}^N H_1^-(\Sigma_{\Delta_i}, \mathbb{Z})//G \qquad \text{(modulo } K\text{-torsion)}\,, \tag{1.1}$$

where G is a certain isotropic subgroup of gluing cycles. Setting up all the right structure and definitions needed to understand (1.1) and the non-abelianization map that relates (1.1) to a statement about gluing equations is a little tricky. In this paper, we will provide the necessary definitions in the case $K = 2$, *i.e.* for spaces of framed flat $PGL(2, \mathbb{C})$ connections. Once the definitions are in place, all proofs are elementary.

We now describe our constructions and main results in a little more detail.

1.1 Symplectic structures from homology

To generalize the notion of an ideal triangulation, we work with a class of "framed" 3-manifolds (Section 2.3). They are oriented manifolds M that can be constructed by gluing together pairs of big, hexagonal faces of truncated tetrahedra (Figure 1). We call this a triangulation $\mathbf{t}$ of M. A framed 3-manifold has its boundary $\mathcal{C} = \partial M$ split into several parts $\mathcal{C} = \mathcal{C}_{\text{big}} \cup \mathcal{C}_{\text{small}} \cup \mathcal{C}_{\text{def}}$. The "big" boundary $\mathcal{C}_{\text{big}}$ is tiled by unglued hexagonal tetrahedron faces (we call this tiling a 2d ideal triangulation $\mathbf{t}_{2d}$), while the "small" boundary $\mathcal{C}_{\text{small}}$ is tiled by the small triangular faces of truncated tetrahedra. If only interiors of some tetrahedron faces are glued there may be also be a "defect" boundary $\mathcal{C}_{\text{def}}$, consisting of annuli around unglued edges.

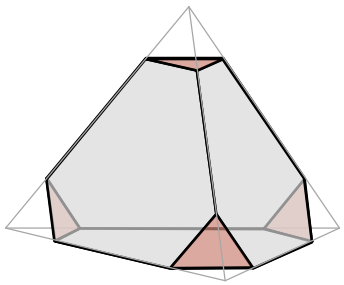

Figure 1. Truncated tetrahedron

One example of a framed 3-manifold is the tetrahedron Δ itself. Its big boundary is a 4-holed sphere and it small boundary contains four discs that fill in the holes. Another example is a cusped hyperbolic manifold, such as a knot complement $M = S^3\backslash K$. Its small boundary consists of a torus T^2 at each cusp, and its big boundary is empty. (An ideal hyperbolic triangulation of M induces a triangulation $\mathbf{t}$ as a framed 3-manifold, with ∂M tiled by truncated vertices of tetrahedra.) A closed hyperbolic 3-manifold with a spun triangulation [11] (*cf.* [43]) is a framed 3-manifold whose boundary only contains small spheres, at the vertices of the spun triangulation. Taking either the cusped or closed hyperbolic examples and deleting all (big) edges of the triangulation $\mathbf{t}$ produces framed 3-manifolds with with defects, $\mathcal{C}_{\text{def}} \neq \oslash$. See also Figures 5–6 on page 14.

Given a framed 3-manifold M, there exists a canonical two-fold branched cover of its boundary $\Sigma \xrightarrow{\pi} \mathcal{C}$. The cover can be constructed by placing a branch point in every face of a triangulation $\mathbf{t}_{2d}$ of $\mathcal{C}_{\text{big}}$, and branch cuts along a trivalent graph dual to the triangulation $\mathbf{t}_{2d}$ (Figure 8, page 15), as well as along the noncontractible cycles of $\mathcal{C}_{\text{def}}$. It turns out that the topological type of the cover is independent of the choice of triangulation used to define it (Lemma 1).

We can define odd homology of the cover as $H_1^-(\Sigma) := \ker\left[\pi_* : H_1(\Sigma) \to H_1(\mathcal{C})\right]$, working implicitly with $\mathbb{Z}$ coefficients. The odd homology is a nondegenerate symplectic lattice, with skew-symmetric product $\langle *, * \rangle$ given by the usual intersection form. If M is such that $\mathcal{C}_{\rm small}$ contains only discs and annuli (say) and $\mathcal{C}_{\rm def}$ is empty, then a short calculation shows that

$$\operatorname{rank} H_1^-(\Sigma) = 6\,\mathrm{genus}(\mathcal{C}) - 6 + 2\,(\#\text{ small discs in } \mathcal{C}_{\rm small})\,. \tag{1.2}$$

(See (2.10) for a more general formula.) We also introduce the twisted homology group $\widetilde{H}_1^-(\Sigma)$, a $\mathbb{Z}_2$ extension of $H_1^-(\Sigma)$, which is defined as the (odd) homology of the unit tangent bundle $T_1\Sigma$ with a $\mathbb{Z}_2$ reduction of the fiber class u (Section 2.2). For the boundary of a framed 3-manifold without defects, there is a natural splitting $\widetilde{H}_1^-(\Sigma) \simeq H_1^-(\Sigma) \oplus \mathbb{Z}_2$ (Lemma 3). It is induced by a surjective map

$$\widetilde{h} : \mathbf{P} \twoheadrightarrow \widetilde{H}_1^-(\Sigma) \tag{1.3}$$

from a certain group $\mathbf{P}$ of paths on $\mathcal{C}_{\rm small}$ (Section 3), whose image is a copy of $H_1^-(\Sigma)$.

One may recognize (1.2) as dimension of Teichmüller space of a punctured surface $\mathcal{C}^*$, formed by puncturing $\mathcal{C}$ once on each small disc. More relevantly for us, it is the complex dimension of the space $\mathcal{X}[\mathcal{C}]$ of framed flat $PGL(2,\mathbb{C})$ connections on $\mathcal{C}^*$, with unipotent holonomy around the punctures. This space is defined fully in Section 4, following [31, 32, 44]; the "framing" of a flat connection consists of an extra choice of invariant flag (*i.e.* an eigenline of the $PGL(2,\mathbb{C})$ holonomy) on every component of $\mathcal{C}_{\rm small}$. We in fact show (Propositions 2–3):

Suppose $\pi_1(\mathcal{C}_{\rm small})$ is abelian. Given any triangulation $\mathbf{t}_{\rm 2d}$ of $\mathcal{C}_{\rm big}$, there is an algebraically open subset $\mathcal{P}[\mathcal{C};\mathbf{t}_{\rm 2d}] \subset \mathcal{X}[\mathcal{C}]$ and a map

$$x : \mathcal{P}[\mathcal{C};\mathbf{t}_{\rm 2d}] \times H_1^-(\Sigma) \to \mathbb{C}^*\,, \tag{1.4}$$

*that's a homomorphism on the second factor (*i.e. $x_{\gamma+\gamma'} = x_\gamma x_{\gamma'}$ *for $\gamma, \gamma' \in H_1^-(\Sigma)$) and nondegenerate in the sense that any basis $\{\gamma_i\}_{i=1}^d$ of $H_1^-(\Sigma)$ provides global coordinates $(x_{\gamma_i}) \in (\mathbb{C}^*)^d$ on $\mathcal{P}[\mathcal{C};\mathbf{t}_{\rm 2d}]$. The map (1.4) may be extended to twisted homology $\widetilde{H}_1^-(\Sigma)$, with the convention that the fiber class u maps to $x_u \equiv -1$. Moreover, there is a non-degenerate holomorphic symplectic structure on $\mathcal{P}[\mathcal{C}]$, which agrees with the Atiyah-Bott structure on the space of ordinary (un-framed) flat connections, whose Poisson brackets are*

$$\{x_\gamma, x'_\gamma\} = \langle \gamma, \gamma' \rangle x_\gamma x_{\gamma'}\,. \tag{1.5}$$

One proof of these statements follows by labeling both cycles $\gamma \in H_1^-(\Sigma)$ and coordinates on $\mathcal{P}[\mathcal{C}]$ by paths $\mathfrak{p} \in \mathbf{P}$, using (1.3), and simply computing Poisson brackets. We will follow this approach in Section 4. More fundamentally, the statements follow from a non-abelianization map

$$\Phi[\mathbf{t}_{\rm 2d}] : \widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma]\big|_R \to \mathcal{P}[\mathcal{C};\mathbf{t}_{\rm 2d}]\,, \tag{1.6}$$

defined using spectral networks in Section 5. Here $\widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma]$ is a moduli space of (twisted and projectivized) flat $GL(1,\mathbb{C})$ connections on Σ, *i.e.* flat $GL(1,\mathbb{C})$ connections on the unit

tangent bundle $T_1\Sigma$ with fiber holonomy -1, modulo a certain projective identification. The space $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ is parametrized by the holonomies of flat connections along cycles $\gamma \in \widetilde{H}^-_1(\Sigma)$, giving an obvious nondegenerate homomorphism

$$x : \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma] \times \widetilde{H}^-_1(\Sigma) \to \mathbb{C}^* , \tag{1.7}$$

and the Atiyah-Bott Poisson bracket among functions x_γ is given trivially by (1.5). We show (Proposition 4) that, subject to some mild restrictions 'R' on the domain, the non-abelianization map $\Phi[\mathbf{t}_{2\rm d}]$ is 1-1 and a symplectomorphism. Therefore, $\mathcal{P}[\mathcal{C}, \mathbf{t}_{2\rm d}]$ inherits the coordinates x_γ and their simple Poisson bracket.

The holomorphic symplectic form ω on $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ has an avatar $\hat{\omega}$ in the K-theory group $K_2(F^*)\otimes\mathbb{Q}$, where F is the field of functions on $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$. It can be written $\hat{\omega} = \frac{1}{2}\sum_{ij}(\epsilon^{-1})^{ij} x_i \wedge x_j$ where $\{x_i\}$ are coordinates associated to a basis $\{\gamma_i\}$ of $H^-_1(\Sigma)$ and $\langle \gamma_i, \gamma_j\rangle =: \epsilon_{ij}$. It follows from the fact that $\Phi[\mathbf{t}_{2\rm d}]$ preserves x_γ functions that the non-abelianization map is in fact a K_2 symplectomorphism, inducing a K_2 avatar of the holomorphic symplectic form on $\mathcal{P}[\mathcal{C}; \mathbf{t}_{2\rm d}]$. Such avatars (and their motivic versions) were first introduced in [44].

1.2 Gluing

Now, suppose that a framed 3-manifold M' is glued together by identifying pairs of hexagonal faces in the big-boundary triangulation $\mathbf{t}_{2\rm d}$ of a framed 3-manifold M. For example, M' could be a knot complement, and M could be a disjoint collection of truncated tetrahedra. We assume (largely for simplicity) that neither M nor M' have defects.

We separate the gluing procedure into two steps. First, by gluing only the interiors of pairs of faces of M we form a framed 3-manifold M_0 that *does* have defects along some edges of its triangulation. Then we fill in the defects to recover M',

$$M \overset{\text{glue interiors of faces}}{\rightsquigarrow} M_0 \overset{\text{fill in edges}}{\rightsquigarrow} M' . \tag{1.8}$$

Each step of the gluing is compatible with the canonical covers Σ, Σ_0, Σ' of the respective boundaries $\mathcal{C}$, $\mathcal{C}_0$, $\mathcal{C}'$. It is then an easy exercise to show that, up to 2-torsion, the homology $\widetilde{H}^-_1(\Sigma')$ is a lattice symplectic reduction of $\widetilde{H}^-_1(\Sigma)$. More precisely (Proposition 1), there is an injection of finite (2-torsion) cokernel

$$\widetilde{g} : \widetilde{H}^-_1(\Sigma_0) \hookrightarrow \widetilde{H}^-_1(\Sigma) , \tag{1.9a}$$

which preserves the intersection form, and there is a distinguished subgroup $\widetilde{G} \subset \widetilde{H}^-_1(\Sigma_0)$ of "gluing cycles" (cycles killed by filling in the defects) such that $\langle \widetilde{G}, \widetilde{G}\rangle = 0$ (*i.e.* $\widetilde{G}$ is isotropic) fitting into the exact sequence (Lemma 3)

$$0 \to \widetilde{G} \xrightarrow{\tilde{i}} \widetilde{K} \xrightarrow{\tilde{q}} \widetilde{H}^-_1(\Sigma') \to 0 , \tag{1.9b}$$

where $\widetilde{K}$ is a finite-index subgroup of the complement $\widetilde{K} \subset \widetilde{K}' := \ker \langle \widetilde{G}, * \rangle \big|_{\widetilde{H}^-_1(\Sigma_0)}$, with $\widetilde{K}'/\widetilde{K}$ = 2-torsion. Thus $\widetilde{H}^-_1(\Sigma') = \widetilde{K}/\widetilde{G} \simeq \widetilde{H}^-_1(\Sigma_0)//\widetilde{G}$ (modulo 2-torsion), and more

generally there's a finite-index sublattice $\widetilde{H} \subset \widetilde{H}_1^-(\Sigma)$ such that

$$\widetilde{H}_1^-(\Sigma') = \widetilde{K}/\widetilde{G} \simeq \widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G}) = \widetilde{H}/\!/g(\widetilde{G})\,. \tag{1.10}$$

(These maps and equivalences hold for un-twisted homology as well.)

The gluing $M \rightsquigarrow M_0 \rightsquigarrow M'$ also induces a gluing of $PGL(2,\mathbb{C})$ and $GL(1,\mathbb{C})$ moduli spaces (Sections 4.4 and 5.2)

$$\begin{aligned} g_{PGL(2)}:\ & \mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]\big|_{x_{g(\tilde{G})}=1,\,R} \twoheadrightarrow \mathcal{P}[\mathcal{C}';\mathbf{t}'_{2\mathrm{d}}]\,, \\ g_{GL(1)}:\ & \widetilde{\mathcal{X}}^-_{\mathrm{ab}}[\Sigma]\big|_{x_{g(\tilde{G})}=1} \twoheadrightarrow \widetilde{\mathcal{X}}^-_{\mathrm{ab}}[\Sigma']\,, \end{aligned} \tag{1.11}$$

where 'R' denotes some mild (open) extra restrictions. In both cases, the gluing maps turn out to be controlled by the gluing equations

$$\boxed{x_{\tilde{g}(\gamma)} = x_{\tilde{q}(\gamma)} \qquad \forall\, \gamma \in \widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)}\,. \tag{1.12}$$

Here the RHS contains all coordinates on $\mathcal{P}[\mathcal{C}']$ (say), since $\widetilde{q}$ is surjective; and they are identified with functions on $\mathcal{P}[\mathcal{C}]$ on the LHS. If $\mu \in \widetilde{G} \subset \widetilde{K}$, then $\widetilde{q}(\mu) = 0$, so the gluing equations simply say $x_{\tilde{g}(\mu)} = 1$, matching the restriction on the domain in (1.11). Our main results then follow quite quickly:

Theorem 1 (p. 41) *The $PGL(2)$ gluing map is the symplectic reduction of a finite quotient*

$$g_{PGL(2)}\ :\ (\mathcal{P}[\mathcal{C}]\big|_R/Z)/\!/(\mathbb{C}^*)^{\mathrm{rank}(\tilde{G})} = (\mathcal{P}[\mathcal{C}]\big|_R/Z)\big|_{\tilde{g}(\tilde{G})=1}/(\mathbb{C}^*)^{\mathrm{rank}(\tilde{G})} \xrightarrow{\sim} \mathcal{P}[\mathcal{C}']\,, \tag{1.13}$$

where 'R' is a mild (open) restriction, $Z \simeq \widetilde{H}_1^-(\Sigma)/\widetilde{H}$ is a finite group (at most 4-torsion), and the group action $(\mathbb{C}^*)^{\mathrm{rank}(\tilde{G})}$ for symplectic reduction is generated by using $x_{\tilde{g}(\mu)}$ ($\mu \in \tilde{G}$) as moment maps with respect to the holomorphic symplectic structure. More so, $g_{PGL(2)}$ is a K_2 symplectic reduction with respect to K_2 avatars $\hat{\omega}$, $\hat{\omega}'$ of the symplectic forms on $\mathcal{P}[\mathcal{C}]$, $\mathcal{P}[\mathcal{C}']$.

An analogue of Theorem 1 holds (rather trivially) for the $GL(1)$ gluing map, and and in fact induces the $PGL(2)$ result via non-abelianization:

Theorem 2 (p. 52) *Gluing and non-abelianization maps fit into a commutative diagram*

$$\begin{array}{ccc} \widetilde{\mathcal{X}}^-_{\mathrm{ab}}[\Sigma]\big|_{x_{\tilde{g}(\tilde{G})=1}} & \overset{g_{GL(1)}}{\twoheadrightarrow} & \widetilde{\mathcal{X}}^-_{\mathrm{ab}}[\Sigma'] = (\widetilde{\mathcal{X}}^-_{\mathrm{ab}}[\Sigma]/Z)/\!/(\mathbb{C}^*)^{\mathrm{rank}\,\tilde{G}} \\ \Phi[\mathbf{t}_{2\mathrm{d}}]\downarrow & & \Phi[\mathbf{t}'_{2\mathrm{d}}]\downarrow \\ \mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]\big|_{x_{\tilde{g}(\tilde{G})=1}} & \overset{g_{PGL(2)}}{\twoheadrightarrow} & \mathcal{P}[\mathcal{C}';\mathbf{t}'_{2\mathrm{d}}] = (\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]/Z)/\!/(\mathbb{C}^*)^{\mathrm{rank}\,\tilde{G}}\,, \end{array} \tag{1.14}$$

with appropriate open restrictions on the domains (as in (1.6), (1.11)); the vertical maps are 1-1 K_2 symplectomorphisms and the horizontal maps are K_2 symplectic reduction.

1.3 Example: knot complement

To finish the introduction, we illustrate in some detail how the abstract formalism described above applies to a simple example, the ideal triangulation of a knot complement. Let $M' = S^3\backslash K$ be the knot complement, viewed as a framed 3-manifold with small torus boundary $\mathcal{C}' = \mathcal{C}'_{\text{small}} \simeq T^2$, and let $M = \sqcup_{i=1}^N \Delta_i$ be the disjoint union of truncated tetrahedra from which M' is glued. Topologically its boundary is a union of spheres, $\mathcal{C} = \sqcup_{i=1}^N \mathcal{C}_{\Delta i}$, $\mathcal{C}_{\Delta i} \simeq S^2$.

The canonical cover of the boundary of a tetrahedron $\Sigma_\Delta \xrightarrow{\pi} \mathcal{C}_\Delta$ is branched over four points, one on each face, as illustrated in Figure 2. Thus it has the topology of a torus, $\Sigma_\Delta \simeq T^2$. The first homology (twisted or untwisted) is entirely odd. A convenient basis of generators is given by cycles γ, γ', γ' encircling pairs of branch points, and thus naturally associated to edges of the tetrahedron. We draw γ in Figure 2, with the convention that solid paths lie on the top sheet of the cover Σ_Δ and dashed lines lie on the bottom sheet. Such smooth curves also have canonical lifts to the unit tangent bundle $T_1\Sigma_\Delta$ (given by their tangent vectors), and thus represent cycles in twisted homology as well. The cycles $\gamma, \gamma', \gamma''$ are equal on opposite edges and sum to zero. Therefore,

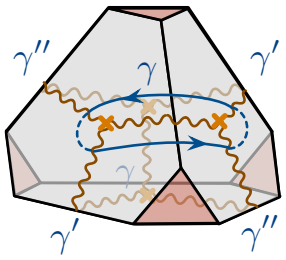

Figure 2. Six edge cycles generating $\widetilde{H}_1^-(\Sigma_\Delta)$.

$$\begin{aligned} H_1^-(\Sigma_\Delta) &= \langle \gamma, \gamma', \gamma'' \,|\, \gamma + \gamma' + \gamma'' = -1 \rangle \simeq \mathbb{Z}^2 \,, \\ \widetilde{H}_1^-(\Sigma_\Delta) &= \langle \gamma, \gamma', \gamma'', u \,|\, \gamma + \gamma' + \gamma'' = -1,\, 2u = 0 \rangle \simeq H_1^-(\Sigma_\Delta) \oplus \mathbb{Z}_2 \,, \end{aligned} \tag{1.15}$$

where u generates the extra fiber class in twisted homology. The intersection product is

$$\langle \gamma, \gamma' \rangle = \langle \gamma', \gamma'' \rangle = \langle \gamma'', \gamma \rangle = 1 \,, \qquad \langle u, * \rangle = 0 \,. \tag{1.16}$$

Correspondingly, the space of framed flat $PGL(2)$ connections on the boundary of a tetrahedron (*cf.* [17, 31]) is

$$\mathcal{P}[\mathcal{C}_\Delta] = \{ x_\gamma, x_{\gamma'}, x_{\gamma''} \,|\, x_\gamma x_{\gamma'} x_{\gamma''} = 1 \} \simeq (\mathbb{C}^*)^2 \,, \tag{1.17}$$

with holomorphic symplectic form, K_2 form, and Poisson brackets

$$\omega = \frac{dx_\gamma}{x_\gamma} \wedge \frac{dx_{\gamma'}}{x_{\gamma'}} \,, \quad \hat{\omega} = x_\gamma \wedge x_{\gamma'} \,; \qquad \{\log x_\gamma, \log x_{\gamma'}\} = \{\log x_{\gamma'}, \log x_{\gamma''}\} = \{\log x_{\gamma''}, \log x_\gamma\} = 1 \,. \tag{1.18}$$

The coordinates on $\mathcal{P}[\mathcal{C}_\Delta]$ (which coincide in this case with Fock-Goncharov coordinates [44], or a complexification of Thurston's shear coordinates) are labelled, as promised, by cycles γ; and their Poisson brackets are induced by the intersection form on $\widetilde{H}_1^-(\Sigma)$.

To connect with hyperbolic geometry, we note that standard shape parameters are related to x coordinates as

$$z = -x_\gamma = x_{\gamma+u} \,, \qquad z' = x_{\gamma'+u} \,, \qquad z'' = x_{\gamma''+u} \,, \tag{1.19}$$

so that $zz'z'' = x_u = -1$. Of course, there is a second standard relation among the shapes in hyperbolic geometry, namely $z + z'^{-1} - 1 = 0$. This second relation describes framed flat connections that extend from the boundary to the *interior* of a tetrahedron, and cuts out a K_2 Lagrangian submanifold $\mathcal{L}_\Delta \subset \mathcal{P}[\mathcal{C}_\Delta]$. We will not need this second relation.

For the full collection M of tetrahedra, we have $H_1^-(\Sigma) = \oplus_{i=1}^N H_1^-(\Sigma_{\Delta i}) \simeq \mathbb{Z}^{2N}$ and $\widetilde{H}_1^-(\Sigma) \simeq H_1^-(\Sigma) \oplus \mathbb{Z}_2$ (in twisted homology, fiber classes on disconnected components are all identified, hence a single $\mathbb{Z}_2$ extension). Correspondingly, $\mathcal{P}[\mathcal{C}] = \prod_{i=1}^N \mathcal{P}[\mathcal{C}_{\Delta i}] \simeq (\mathbb{C}^*)^{2N}$. We denote the generators of $\widetilde{H}_1^-(\Sigma)$ as $\gamma_i, \gamma_i', \gamma_i''$ and the associated functions on $\mathcal{P}[\mathcal{C}]$ as $z_i = -x_{\gamma i}$, $z_i' = -x_{\gamma i'}$, etc., with index i for the i-th tetrahedron.

For the glued-up knot complement M', the canonical cover of the boundary $\Sigma' \xrightarrow{\pi} \mathcal{C}' \simeq T^2$ is unbranched and disconnected, *i.e.* $\Sigma' \simeq T^2 \sqcup T^2$. Therefore,

$$\begin{aligned} H_1^-(\Sigma') &= \mathbb{Z}\langle \alpha, \beta \rangle \simeq \mathbb{Z}^2 \,, \\ \widetilde{H}_1^-(\Sigma') &= \mathbb{Z}\langle \alpha, \beta, u \,|\, 2u = 0 \rangle = H_1^-(\Sigma') \oplus \mathbb{Z}_2 \,, \end{aligned} \tag{1.20}$$

where α and β are "odd double lifts" of (say) the meridian and longitude cycles on the torus boundary of M', as on the left of Figure 3. The intersection product is

$$\langle \alpha, \beta \rangle = 2 \,, \tag{1.21}$$

with $\langle u, * \rangle \equiv 0$ (always). Correspondingly, there is a complex symplectic space [12, 21]

$$\mathcal{P}[\mathcal{C}'] = \{x_\alpha, x_\beta\} \simeq (\mathbb{C}^*)^2 \,, \qquad x_\alpha = \ell^2 \,,\ x_\beta = m^2 \,, \tag{1.22}$$

of framed flat $PGL(2)$ connections on the torus boundary of the knot complement, parametrized by the squares of meridian and longitude eigenvalues $x_\alpha = \ell^2$ and $x_\beta = m^2$. Now

$$\omega' = \frac{1}{2} \frac{dx_\alpha}{x_\alpha} \wedge \frac{dx_\beta}{x_\beta} \,, \quad \hat{\omega}' = \frac{1}{2} x_\alpha \wedge x_\beta \,; \qquad \{\log x_\alpha, \log x_\beta\} = 2 \,. \tag{1.23}$$

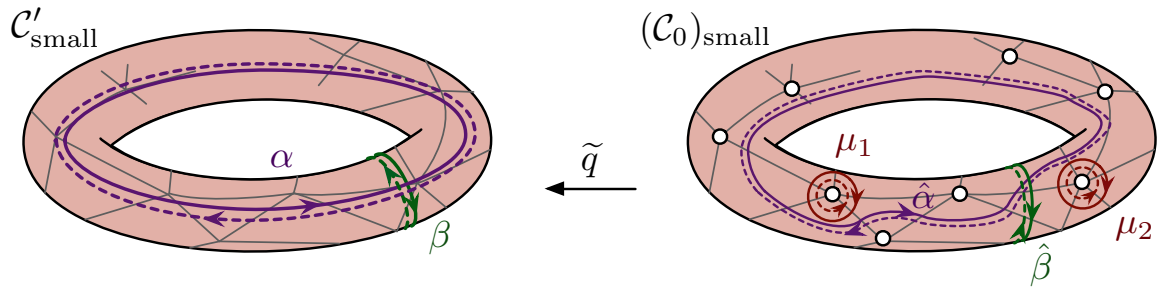

Figure 3. The small boundaries of M' and M_0, with their (unbranched) canonical covers Σ', Σ_0, and curves representing generators of $\widetilde{H}_1^-(\Sigma')$ and $\widetilde{H}_1^-(\Sigma_0)$.

In order to relate the boundaries of the tetrahedra M and the knot complement M', we form an intermediate framed 3-manifold M_0 with defects: M_0 is obtained by gluing tetrahedra

(only) along the interiors of their big faces, or equivalently by removing the edges from the triangulation $\mathbf{t}$ of M'. If there are N tetrahedra there are also N edges in the triangulation. Therefore, the boundary $\mathcal{C}_0 = \partial M_0$ is topologically a surface of genus $N+1$. Its small part $(\mathcal{C}_0)_{\text{small}}$ is a torus with $2N$ holes (formed from the torus $\mathcal{C}'_{\text{small}}$ with endpoints of edges removed), and its defect part $(\mathcal{C}_0)_{\text{def}}$ consists of N annuli. (In this case $(\mathcal{C}_0)_{\text{big}} = \oslash$.) The canonical cover $\Sigma_0 \xrightarrow{\pi} \mathcal{C}_0$ has no branch points, but is characterized by a branch cut along the non-contractible cycle of each annulus in $(\mathcal{C}_0)_{\text{def}}$. It has genus $2N+1$, whence

$$\operatorname{rank} H_1^-(\Sigma_0) = \operatorname{rank} H_1(\Sigma_0) - \operatorname{rank} H_1(\mathcal{C}_0) = 2(2N+1) - 2(N+1) = 2N\,; \tag{1.24}$$

thus $H_1^-(\Sigma_0) \simeq \mathbb{Z}^{2N}$ and $\widetilde{H}_1^-(\Sigma_0)$ is a $\mathbb{Z}_2$ extension thereof. The finite-cokernel injection $\widetilde{g}$ in (1.9) maps $\widetilde{H}_1(\Sigma_0)$ into $\widetilde{H}_1(\Sigma) \simeq \mathbb{Z}^{2N} \oplus \mathbb{Z}_2$, preserving the fiber class u.

Inside $\widetilde{H}_1^-(\Sigma_0)$ lies the subgroup $\widetilde{G}$ of gluing cycles, generated by "odd double lifts" μ_j of curves surrounding endpoints of defects I_j on $(\mathcal{C}_0)_{\text{small}}$, as on the right of Figure 3. (The two cycles at the two endpoints of a defect I_j are equivalent, so they may unambiguously be called μ_j.) It is easy to see that the only relation among the μ_j is $\sum_{j=1}^N \mu_j = 0$, whence

$$\widetilde{G} \simeq \mathbb{Z}^{N-1}\,. \tag{1.25}$$

Moreover, since $\langle \mu_j, \mu_{j'} \rangle = 0$, $\widetilde{G}$ is an isotropic subgroup of $\widetilde{H}_1^-(\Sigma_0)$. The "complement" $\widetilde{K} := \ker \langle \widetilde{G}, * \rangle|_{\widetilde{H}_1^-(\Sigma_0)}$ includes $\widetilde{G}$ itself and is a subgroup of rank $N+1$,

$$\widetilde{G} \subset \widetilde{K} \simeq \mathbb{Z}^{N+1} \oplus \mathbb{Z}_2\,. \tag{1.26}$$

The two additional generators of $\widetilde{K}$ (besides the fiber class u) are represented by "odd double lifts" $\hat{\alpha}, \hat{\beta}$ of any curves on $(\mathcal{C}_0)_{\text{small}}$ that map to α, β once the holes in $(\mathcal{C}_0)_{\text{small}}$ are filled in. The map that fills in the holes is $\widetilde{q}$ from (1.9b); it sends $(\hat{\alpha}, \hat{\beta}, \mu_j)$ to $(\alpha, \beta, 0)$. Thus

$$\widetilde{H}_1^-(\Sigma') \simeq \widetilde{K}/\widetilde{G} \simeq \widetilde{H}_1^-(\Sigma_0)//\widetilde{G}\,. \tag{1.27}$$

Now, Thurston's gluing equations state that the product of shapes z_i, z_i', z_i'' around any edge I_j in the triangulation of M' is trivial, and that the product around longitude and meridian paths (with exponents ± 1) equals the squares ℓ^2, m^2 of longitude and meridian eigenvalues (*cf.* [12]). In terms of homology, it turns out that the injection $\widetilde{g}$ from (1.9a) also sends each μ_j to a sum of tetrahedron cycles $\gamma_i, \gamma_i', \gamma_i''$ around edge I_j, and sends $\hat{\alpha}, \hat{\beta}$ to sums and differences of parameters around longitude and meridian paths. We will show in Section 4.6 that, together with some signs coming from the fiber class u, Thurston's gluing equations are precisely and succinctly written as

$$x_{\tilde{q}(\gamma)} = x_{\tilde{g}(\gamma)}\,, \qquad \forall\, \gamma \in \widetilde{K}\,. \tag{1.28}$$

When $\gamma = \mu_j$ for some j, this is an edge equation; and when $\gamma = \hat{\alpha}$ or $\hat{\beta}$ these are cusp equations. It follows (Theorem 1) that $\mathcal{P}[\mathcal{C}'] = (\mathcal{P}[\mathcal{C}]/Z)//(\mathbb{C}^*)^{N-1} = (\mathcal{P}[\mathcal{C}]/Z)|_{\tilde{g}(\mu_j)=1}/(\mathbb{C}^*)^{N-1}$ is the symplectic reduction of a finite quotient. Explicitly, the quotient action is $x_\gamma \sim$

$(t^\mu)^{\langle\gamma,\tilde{g}(\mu)\rangle}x_\gamma$ for all $\gamma \in \widetilde{H}_1^-(\Sigma)$ and $\mu \in \widetilde{G}$, where (say) if $\mu = \sum_{j=1}^{N-1} a_j\mu_j$ in any basis for $\widetilde{G}$ then $t^\mu \in \mathbb{C}^*$ is defined as $t_1^{a_1}\cdots t_{N-1}^{a_{N-1}}$ ($t_j \in \mathbb{C}^*$).

There is also a slightly weaker algebraic version of the symplectic properties. Choosing $\{\gamma_i, \gamma_i'\}_{i=1}^N$ as a basis for $H_1^-(\Sigma)$ and $z_i = -x_{\gamma i}$, $z_i' = -x_{\gamma i'}$ as the corresponding coordinates on $\mathcal{P}[\mathcal{C}]$, the gluing equations take the form

$$\begin{aligned} \ell^2 = x_\alpha \quad &= \pm z^A z'^{A'} \\ m^2 = x_\beta \quad &= \pm z^B z'^{B'} \\ 1 = x_{\tilde{q}(\mu_j)} \quad &= \pm z^{C_j} z'^{C_j'} \qquad (\forall\ 1 \le j \le N) \end{aligned} \tag{1.29}$$

for N-dimensional integer vectors A, A', B, B' and $N \times N$ matrices C, C'. The $(N+1) \times 2N$ matrix

$$\mathbf{g} = \begin{pmatrix} A & A' \\ B & B' \\ C & C' \end{pmatrix} \tag{1.30}$$

is nothing but the matrix of the injection $\widetilde{g}$ restricted to $\widetilde{K}$ (ignoring the fiber class u). The fact that $\widetilde{g}$ preserves the intersection form implies

$$\mathbf{g}\, J_{2N}\, \mathbf{g}^T = (2J_2) \oplus 0_{N\times N}\,, \qquad J_{2n} := \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}\,; \tag{1.31}$$

while because $\operatorname{rank}(\widetilde{K}) = N+1$ we must have $\operatorname{rank}(\mathbf{g}) = N+1$. This was precisely the form of the symplectic properties initially discovered by Neumann and Zagier [12].

2 Preliminaries

We begin by defining the odd homology of a cover and its twisted version. Then we formally introduce the "framed" 3-manifolds that support framed flat connections and construct canonical covers of their boundaries. Section 2.5 contains the basic but central result that gluing of framed 3-manifolds induces a lattice symplectic reduction on the odd homology of their boundaries.

2.1 Odd homology

Let $\mathcal{C}$ be a closed, oriented surface or disjoint union thereof, and let $\pi : \Sigma \to \mathcal{C}$ be an oriented double cover, branched over a finite (possibly empty) collection of isolated points $p \in \mathfrak{b}$. We can think of $\mathfrak{b}$, the branching locus, as a subset of either $\mathcal{C}$ or Σ. We always assume that the branching is simple, with ramification index 2 at any $p \in \mathfrak{b}$.

We define the "odd" homology of the cover (with $\mathbb{Z}$ coefficients always assumed) as

$$H_\bullet^-(\Sigma) := \ker\left[\pi_* : H_\bullet(\Sigma) \to H_\bullet(\mathcal{C})\right]. \tag{2.1}$$

Note that, since Σ is oriented, $H_\bullet(\Sigma)$ is torsion-free, whence $H_\bullet^-(\Sigma)$ is torsion-free as well. It is also convenient to introduce the deck transformation $\sigma : \Sigma \to \Sigma$, an orientation-preserving

involution. Letting $\sigma_* : H_\bullet(\Sigma) \xrightarrow{\sim} H_\bullet(\Sigma)$ denote the induced automorphism on homology and $P_\pm := \mathbb{1} \pm \sigma_*$ the associated (quasi-)projections, we may equivalently define

$$H_\bullet^-(\Sigma) := \ker P_+ = \ker \pi_* \,, \qquad H_\bullet^+(\Sigma) := \ker P_- \,. \tag{2.2}$$

The equivalence of (2.1) and (2.2) follows from the existence of an injection $\ell^+ : H_\bullet(\mathcal{C}) \hookrightarrow H_\bullet(\Sigma)$ that sends a cycle to the sum of its pre-images[4], obeying

$$\pi_* \circ \ell^+ = 2\mathbb{1}_{H_\bullet(\mathcal{C})} \,, \qquad \ell^+ \circ \pi_* = P_+ \,. \tag{2.3}$$

We call ℓ^+ an "even double lift." A more careful description of ℓ^+ appears in Appendix A, together with some other simple results about odd homology.

The first homology group $H_1(\Sigma)$ has a non-degenerate skew-symmetric intersection form $\langle\ ,\ \rangle : \bigwedge^2 H_1(\Sigma, \mathbb{Z}) \to \mathbb{Z}$, which is preserved by σ_* (since σ is orientation-preserving). It follows (Lemma 7, Appendix A) that the intersection form is non-degenerate on $H_1^-(\Sigma)$ as well.

It follows from injectivity of ℓ^+ that

$$\operatorname{rank} H_1^-(\Sigma) = \operatorname{rank} H_1(\Sigma) - \operatorname{rank} H_1(\mathcal{C}) \,. \tag{2.4}$$

Combined with the Riemann-Hurwitz formula $\chi(\Sigma) = 2\chi(\mathcal{C}) - \#(\mathfrak{b})$ (see (A.7)), we also obtain

$$\operatorname{rank} H_1^-(\Sigma) = -\chi(\mathcal{C}) + \#(\mathfrak{b}) \,. \tag{2.5}$$

2.2 Twisted homology

We will often need a $\mathbb{Z}_2$ extension of odd homology. Let S be any closed, oriented, connected surface and let T_1S denote its unit tangent bundle. The homology $H_1(T_1S)$ is an extension[5] of $H_1(S)$ by $\mathbb{Z}/(\chi(S)\mathbb{Z})$, where $\mathbb{Z}/(\chi(S)\mathbb{Z})$ is generated by the class u of the unit-tangent fiber above any point of S, which satisfies $\chi(S)\, u = 0$. We define the twisted homology $\widetilde{H}_1(S) := H_1(T_1S)/\mathbb{Z}\langle 2u\rangle$ to be the quotient of $H_1(T_1S)$ by the subgroup generated by $2u$; this makes sense since $\chi(S) = 2 - 2g(S)$ is always even. Thus $\widetilde{H}_1(S)$ is a $\mathbb{Z}_2$ extension,

$$\begin{aligned} 0 \to \quad \mathbb{Z}_2 \quad &\hookrightarrow \widetilde{H}_1(S) \xrightarrow{p} H_1(S) \to 0 \,, \\ = \mathbb{Z}\langle u\rangle/\mathbb{Z}\langle 2u\rangle \quad & \end{aligned} \tag{2.6}$$

where p just sets $u \mapsto 0$. More generally, if $S = \sqcup_{i=1}^n S_i$ has multiple connected components S_i, we define $\widetilde{H}_1(S) := \big[\oplus_{i=1}^n \widetilde{H}_1(S_i)\big]/\mathbb{Z}\langle u_i - u_j\rangle$ by identifying all the fiber classes $u_i \sim u_j$ $(1 \le i, j \le n)$. Then $\widetilde{H}_1(S)$ is still a single $\mathbb{Z}_2$ extension of $H_1(S)$, as in (2.6).

We can represent all classes $[\gamma] \in \widetilde{H}_1^-(S)$ by drawing smooth curves γ on S, and using their unit tangent vectors to lift them canonically to curves on T_1S. In particular, the fiber class u is represented by any small contractible loop, as in Figure 4.

[4] This is sometimes called a "transfer homomorphism," *cf.* [45, Chapter G.3]

[5] Showing this is a classic application of the Gysin sequence.

Figure 4. Two smooth curves on a surface S representing $[\gamma]$ and $[\gamma'] = [\gamma] + u$ in $\widetilde{H}_1(S)$.

Now suppose $\Sigma \xrightarrow{\pi} \mathcal{C}$ is an oriented branched double cover as above. The deck transformation σ has nonvanishing derivative $d\sigma$, and thus extends to an involution on $T_1\Sigma$, which in turn induces an involution σ_* on twisted homology $\widetilde{H}_1(\Sigma)$. Indeed, since the deck transformation is orientation-preserving, we have $\sigma_* u = u$. Letting $P_\pm = \mathbb{1} \pm \sigma_*$, we define

$$\widetilde{H}_1^\pm(\Sigma) := \ker P_\mp\big|_{\widetilde{H}_1(\Sigma)} . \tag{2.7}$$

Note that $P_+ u = P_- u = 0$, so $u \in \widetilde{H}_1^\pm(\Sigma)$. Indeed, $\widetilde{H}_1^-(\Sigma)$ (say) is just the preimage of $H_1^-(\Sigma) \subset H_1(\Sigma)$ under the map p in (2.6), so $\widetilde{H}_1^-(\Sigma)$ is a $\mathbb{Z}_2$ extension of $H_1^-(\Sigma)$,

$$0 \to \mathbb{Z}_2 \hookrightarrow \widetilde{H}_1^-(\Sigma) \xrightarrow{p} H_1^-(\Sigma) \to 0 . \tag{2.8}$$

The intersection product on $\widetilde{H}_1^-(\Sigma)$ is simply pulled back from $H_1^-(\Sigma)$, *i.e.* $\langle \gamma, \gamma' \rangle_{\widetilde{H}_1^-} := \langle p(\gamma), p(\gamma') \rangle_{H_1^-}$. Since the product on $H_1^-(\Sigma)$ is nondegenerate, it follows that $\langle \gamma, \gamma' \rangle = 0$ $\forall\, \gamma' \in \widetilde{H}_1^-(\Sigma)$ if and only if $\gamma = 0$ or $\gamma = u$.

2.3 Boundaries of framed 3-manifolds

Framed 3-manifolds, as introduced in [31, 32], are a class of 3-manifolds with extra structure on which it is natural to define moduli spaces of framed flat connections. They include both knot complements and single tetrahedra. In the present section, we look at framed 3-manifolds topologically, construct canonical branched covers of their boundaries, and study how the homology of these covers behaves under gluing. We use the most liberal possible definition:

Definition 1 *A* framed 3-manifold M *is an oriented 3-manifold with boundary (or disjoint union of such), together with a splitting of its boundary* $\partial M = \partial M_{\text{big}} \cup \partial M_{\text{small}}$ *into "big" and "small" parts, such that*

a) ∂M_{big} *consists of surfaces of any genus with at least one hole and negative Euler character;*

b) $\partial M_{\text{small}}$ *consists of surfaces of any genus with any number of holes;*

c) the two parts of the boundary attach along circles, $\partial M_{\text{big}} \cap \partial M_{\text{small}} = \sqcup_i S_i^1$.

Often (as in [31, 32]) it is useful to impose additional restrictions on the small boundary. For example, requiring $\partial M_{\text{small}}$ to have abelian fundamental group ensures that framing data does not restrict the choice of a flat connection on ∂M; and requiring $\partial M_{\text{small}}$ to admit a Euclidean

structure ensures the existence of additional data needed for quantization (*e.g.* combinatorial flattenings of [29], *cf.* [17, 19]). For the moment, we don't need these restrictions.

The basic example of a framed 3-manifold is a truncated tetrahedron (Figure 1), whose big boundary is a 4-holed sphere (formed from four big hexagonal faces) and whose small boundary is four small triangles (at the truncated vertices). We define a *triangulation* (or tiling) $\mathbf{t}$ of a framed 3-manifold to be a decomposition into truncated tetrahedra that are glued only along big hexagonal faces.[6] Thus, $\mathbf{t}$ induces

- a tiling of ∂M_{big} by big hexagons, equivalent to an ideal 2d triangulation $\mathbf{t}_{2d}$; and
- a tiling of $\partial M_{\text{small}}$ by small triangles.

For example, any knot complement M with a standard ideal triangulation can be viewed as a framed 3-manifold with a single small torus boundary and a tiling $\mathbf{t}$ where all big hexagonal faces have been glued pairwise. The tiling is simply obtained by truncating the ideal triangulation. This is illustrated in Figure 5.

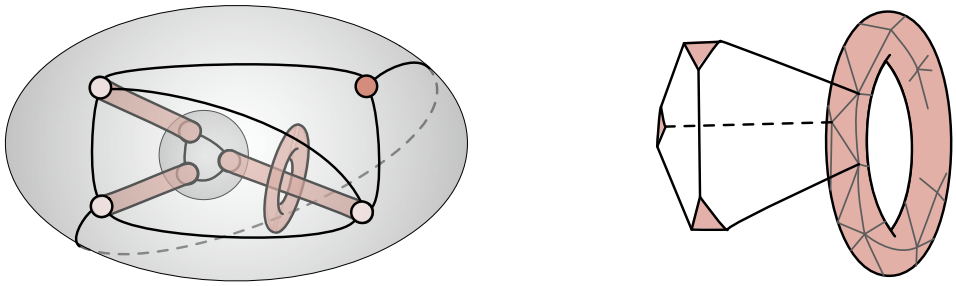

Figure 5. Left: a typical framed 3-manifold, with two big boundaries (a 3-holed sphere and a 4-holed sphere) and several small tori, annuli, and discs. A 2d triangulation of the big boundaries is shown. Right: the triangulation of a small torus boundary induced from a 3d triangulation $\mathbf{t}$.

We also introduce

Definition 2 *A* framed 3-manifold with defects M *is an oriented 3-manifold with boundary, together with a splitting* $\partial M = \partial M_{\text{big}} \cup \partial M_{\text{small}} \cup \partial M_{\text{def}}$ *into three parts, such that* ∂M_{big} *and* $\partial M_{\text{small}}$ *are as in Def. 1, while* ∂M_{def} *is a collection of annuli that attach only to holes on the small boundary:* $\partial M_{\text{big}} \cap \partial M_{\text{small}} = \sqcup_i S^1_i$, $\partial M_{\text{small}} \cap \partial M_{\text{def}} = \sqcup_i S^1_j$, $\partial M_{\text{def}} \cap \partial M_{\text{big}} = \oslash$.

We can think about forming a framed 3-manifold with defects in two different ways. We can start with a framed 3-manifold M without defects and a collection of open curves $\{I_i\}$ in the interior of M, whose endpoints lie on the small boundary. Excising neighborhoods of the I_i produces a framed 3-manifold with defects $M' = M \setminus \cup_i I_i$, as in Figure 6. We say

[6] Such a triangulation always exists, though we don't need this fact here. For example, to construct a triangulation for an arbitrary framed 3-manifold M, one can first take two copies $M, \overline{M}$ with opposite orientation and glue them along the big boundary to form $N = M \cup_{\text{big}} \overline{M}$. The new N has only closed small boundary components. It admits a topological ideal triangulation (equivalent to a tiling by truncated tetrahedra) by classic theorems of Matveev [46] (*cf.* [47, Prop 1.2]). Then one can refine the triangulation of N so that $\partial M_{\text{big}} \subset N$ is realized by faces of tetrahedra, and cut along ∂M_{big} to recover a triangulation of M.

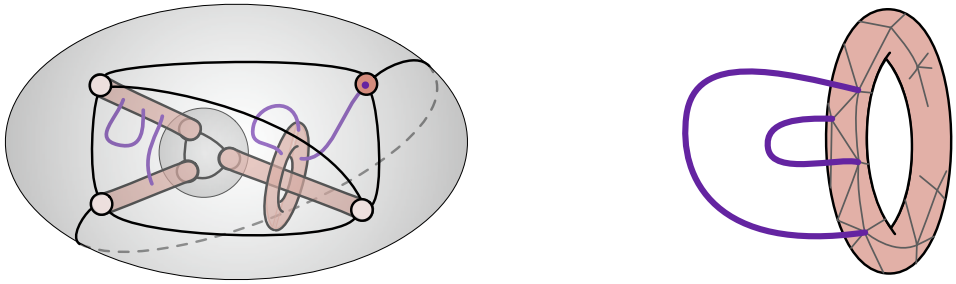

Figure 6. Including defects (extra annular boundary components) that begin and end on $\partial M_{\text{small}}$. On the right, the defects are compatible with a 3d triangulation: they begin and end at vertices of the tiling of $\partial M_{\text{small}}$.

that a triangulation $\mathbf{t}$ of M is compatible with the defects if the I_i lie along edges of the triangulation. Then M' inherits a tiling $\mathbf{t}'$ by tetrahedra with truncated vertices and some truncated edges (those along the I_i), as in Figure 7.

Alternatively, we may start with a framed 3-manifold M and a triangulation $\mathbf{t}$, and identify the *interiors* (but not the edges) of some pairs of faces of the truncated tetrahedra in M. In general this again leads to a framed 3-manifold M' with defects along the new internal edges created by the gluing. It inherits a triangulation $\mathbf{t}'$.

In the extreme case, we may build a framed 3-manifold M' with defects entirely from doubly truncated tetrahedra, by gluing their big hexagonal faces together in pairs. Then a defect lies along every edge in the triangulation of M'.

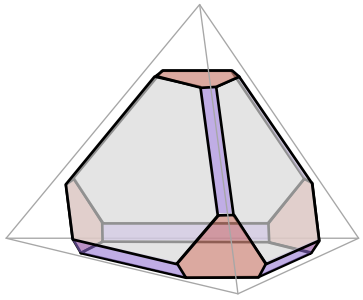

Figure 7. Doubly truncated tetrahedron

2.4 The canonical cover

The boundary $\mathcal{C} = \partial M$ of any framed 3-manifold, with or without defects, admits a canonical double cover $\Sigma \xrightarrow{\pi} \mathcal{C}$.

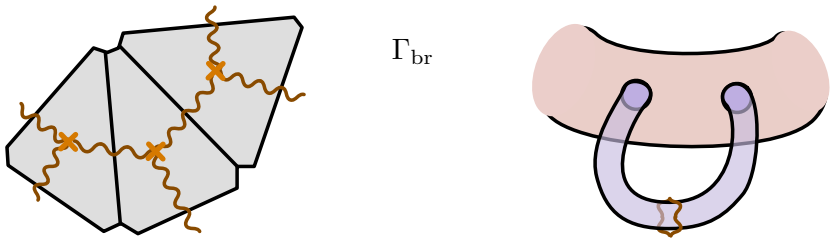

Figure 8. The branching graph Γ_{br}, dual to a big-boundary triangulation (left) and containing closed curves around each defect (right).

To construct it, we first choose a triangulation $\mathbf{t}$ of M. Let $\mathbf{t}_{2\text{d}}$ be the induced tiling of

the big boundary $\partial M_{\rm big}$ by hexagons, thought of as a 2d ideal triangulation. Let $\Gamma_{\rm br}$ be the union of a trivalent graph dual to the triangulation $\mathbf{t}_{\rm 2d}$ and a closed curve around the girth of every defect of M, as in Figure 8. We construct the cover $\Sigma \xrightarrow{\pi} \mathcal{C} = \partial M$ by using $\Gamma_{\rm br}$ as a branching graph — putting branch cuts on the edges of $\Gamma_{\rm br}$ and branch points at the vertices of $\Gamma_{\rm br}$. In other words, we take two identical copies Σ^+, Σ^- of $\mathcal{C}\backslash\Gamma_{\rm br}$, glue Σ^+ to Σ^- (and vice versa) along the edges of $\Gamma_{\rm br}$, and add in one copy of the vertices $\mathfrak{b}$ of $\Gamma_{\rm br}$ to complete the space to a closed surface Σ:

$$\Sigma = \Sigma^+ \cup \Sigma^- \cup \mathfrak{b} = (\mathcal{C}\backslash\Gamma_{\rm br}) \cup_{\rm edges(\Gamma_{br})} (\mathcal{C}\backslash\Gamma_{\rm br}) \cup \mathfrak{b} \qquad (\text{with } \mathfrak{b} = \text{vertices}(\Gamma_{\rm br}))\,. \tag{2.9}$$

Lemma 1 *This definition of the canonical cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ does not depend on a choice of triangulation $\mathbf{t}$.*

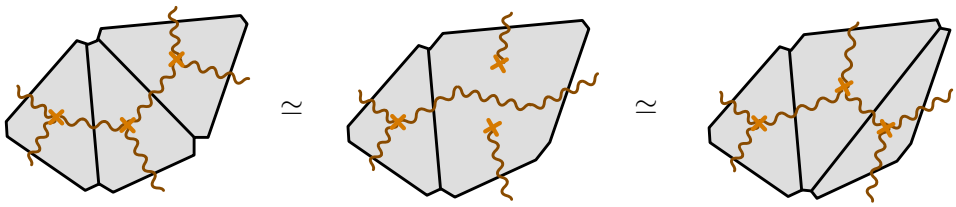

Figure 9. Invariance of the canonical cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ under a flip of the big-boundary triangulation.

Proof. The only possible dependence on $\mathbf{t}$ comes from the induced ideal triangulation $\mathbf{t}_{\rm 2d}$ on the big boundary. Any two triangulations of the big boundary are related by a sequence of flips, *a.k.a.* 2-2 Pachner moves. Each flip, moreover, changes the cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ locally in a way that does not modify its topological type (*i.e.*, corresponds to a homeomorphism of Σ and a homotopy of π). To see this, we simply reconnect branch cuts as in Figure 9. (In general, the topological type of a cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ defined from a branching graph only depends on the relative homology class in $H_1(\mathcal{C}, \mathfrak{b}; \mathbb{Z}_2)$ induced by the graph.) $\square$

Another simple but fundamental result is

Lemma 2 *Let M be a framed 3-manifold without defects, with boundary $\mathcal{C} = \partial M$ such that $\pi_1(\mathcal{C}_{\rm small})$ is abelian. Let $\mathcal{C}_i$ denote the connected components of $\mathcal{C}$, with g_i the genus of $\mathcal{C}_i$ and n_i the number of disc components of $(\mathcal{C}_i)_{\rm small}$. Let $\Sigma_i \xrightarrow{\pi} \mathcal{C}_i$ be the canonical cover over $\mathcal{C}_i$. Then*

$$\operatorname{rank} H_1^-(\Sigma_i) = \begin{cases} 0 & \mathcal{C}_i \text{ is a small sphere} \\ 2 & \mathcal{C}_i \text{ is a small torus} \\ 6g_i - 6 + 2n_i & \text{otherwise}\,. \end{cases} \tag{2.10}$$

Proof. This is trivial for small spheres or tori (which themselves are closed, disjoint components of $\mathcal{C}$). Otherwise, $\mathcal{C}_i$ contains big boundary with holes connected to small discs and small annuli. We calculate $\chi(\mathcal{C}_i) = (\#\text{ faces of } \mathbf{t}_{\rm 2d}) + n_i - (\#\text{ big edges of } \mathbf{t}_{\rm 2d}) =$

$-\frac{1}{2}(\#\text{ faces of } \mathbf{t}_{2d}) + n_i = -\frac{1}{2}\#(\mathfrak{b}) + n_i$. On the other hand, $\chi(\mathcal{C}_i) = 2 - 2g_i$, so $\#(\mathfrak{b}) = 4g_i - 4 + 2n_i$; then the result follows from (2.5). □

2.5 Gluing of canonical covers

As a final basic result, we now describe how the odd homology of canonical covers of boundaries changes as framed 3-manifolds are glued.

Consider the following setup, used throughout the remainder of the paper. Let M be a framed 3-manifold with triangulation $\mathbf{t}$. (M may have defects and/or multiple components.) Choose several pairs of faces (f_i, f_i') of the induced triangulation on the big boundary ∂M_{big} to glue together, and identify their *interiors* to form a manifold M_0:

$$M_0 = M|_{f_i^\circ \sim -f_i'^\circ}\,. \tag{2.11}$$

In general, M_0 will be a framed 3-manifold with defects, and an induced triangulation $\mathbf{t}_0$. Let $\{I_j\}$ denote the new defects in M_0, corresponding to new internal edges of $\mathbf{t}_0$ that were not present in $\mathbf{t}$. We may then fill in the defects I_j to form a new 3-manifold M',

$$M' = M_0 \cup (\sqcup_j I_j)\,. \tag{2.12}$$

We illustrate this gluing process in Figure 10.

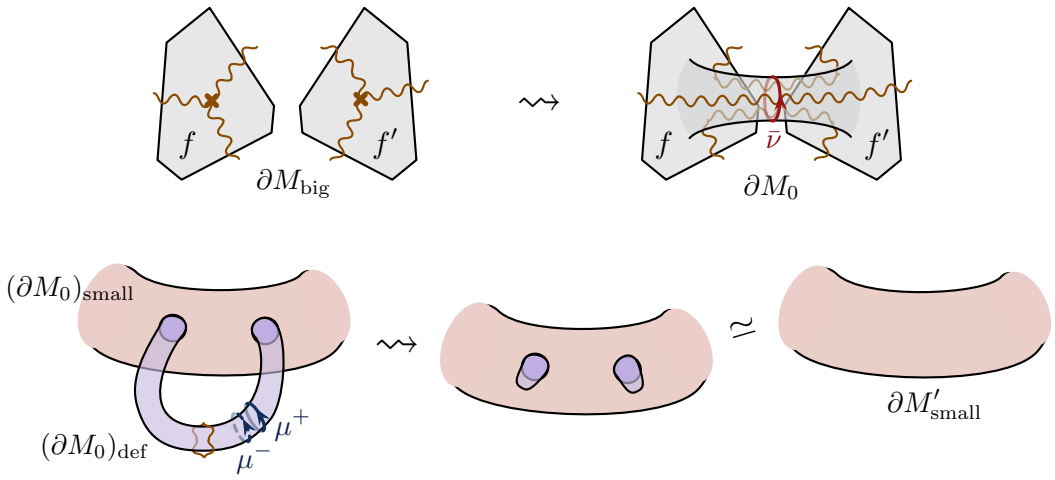

Figure 10. Gluing $M \rightsquigarrow M_0 \rightsquigarrow M'$ in two steps: first, identifying interiors of faces on the big boundary of M, then filling in defects of M_0.

A standard example of such gluing involves a disjoint union of truncated tetrahedra $M = \sqcup_{i=1}^N \Delta_i$, glued together to form (say) a knot complement M'. In this case M_0 is obtained by gluing doubly truncated tetrahedra; the boundary ∂M_0 has genus $N+1$ and can be thought of as the small torus boundary of M' with N additional defects drilled out. We will come back to this example in Sections 3.4, 4.6.

Now, let us denote the boundaries of the three 3-manifolds as $\mathcal{C} := \partial M$, $\mathcal{C}_0 := \partial M_0$ and $\mathcal{C}' := \partial M'$. They have canonical double covers

$$\Sigma \xrightarrow{\pi} \mathcal{C}\,, \qquad \Sigma_0 \xrightarrow{\pi} \mathcal{C}_0\,, \qquad \Sigma' \xrightarrow{\pi} \mathcal{C}'\,. \tag{2.13}$$

To relate the homologies of these covers, we introduce oriented closed curves $\overline{\mu}_j \subset \mathcal{C}_0$ winding around the girth of each defect I_j (parallel to the branch lines, as in Figure 10). Denote by μ_j^+, μ_j^- the lifts of $\overline{\mu}_i$ to the two sheets of Σ_0, and define subgroups

$$\begin{array}{ll} G_0 \subset H_1(\Sigma_0) & \text{subgroup generated by all the } \mu_j^\pm \\ G' := \ker P_+|_{G_0} \subset H_1^-(\Sigma_0) & \\ G := \operatorname{im} P_-|_{G_0} \subset H_1^-(\Sigma_0) & \text{subgroup generated by } \mu_j := \mu_j^+ - \mu_j^- \end{array} \tag{2.14}$$

In general, G'/G may be a nontrivial 2-torsion group.

Proposition 1 *Let $K' := \ker \langle G', * \rangle|_{H_1^-(\Sigma_0)}$. Then:*

a) There is an injection $g : H_1^-(\Sigma_0) \hookrightarrow H_1^-(\Sigma)$ with finite (2-torsion) cokernel that preserves the intersection form.

b) There is an exact sequence

$$0 \to G' \xrightarrow{i} K' \xrightarrow{q} H_1^-(\Sigma') \to 0 \tag{2.15}$$

providing an isomorphism $H_1^-(\Sigma') \simeq H_1^-(\Sigma_0)//G' = K'/G'$, preserving the intersection form.

c) Suppose $K \subset H_1^-(\Sigma_0)$ is a subgroup such $G \subset K \subset K'$ and $K'/G' \simeq K/G$. Then there is a finite-index subgroup $H \subset H_1^-(\Sigma)$ containing $g(G)$ such that

$$H_1^-(\Sigma') \simeq H//g(G) = g(K)/g(G)\,. \tag{2.16}$$

The quotient $H_1^-(\Sigma)/H$ contains (at most) 2-torsion and 4-torsion.

Proof. For part (a), we observe that relating $\Sigma \rightsquigarrow \Sigma_0$ involves the inverse of the basic cut-and-glue operation of Appendix A.4. Namely, from the perspective of the boundary $\mathcal{C}$, gluing the interiors of two faces (f_i, f_i') is equivalent to excising discs in the interiors of these faces (containing branch points) and subsequently identifying the circular boundaries of holes that are created. The gluing can be reversed by cutting the boundary $\mathcal{C}_0$ along curves $\overline{\nu}_i$ (Figure 10), and filling in the resulting holes with discs. This gluing/ungluing lifts to Σ and Σ_0. Applying Lemma 14 (p. 59) to the pair (Σ_0, Σ), and noting that the group $G_\nu \subset H_1(\Sigma_0)$ generated by lifts ν_i of the $\overline{\nu}_i$ is even — so that $G_\nu^- := \ker P_+|_{G_\nu} = 0$ and $\ker \langle G_\nu^-, * \rangle|_{H_1^-(\Sigma_0)} = 0$ — we obtain the injection $H_1^-(\Sigma_0)//G_\nu^- = H_1^-(\Sigma_0) \hookrightarrow H_1^-(\Sigma)$. The cokernel is at most 2-torsion and may be nontrivial because the lifts ν_i are connected (and branch points are destroyed/created in the gluing/ungluing).

For part (b), we observe that relating boundaries $\mathcal{C}_0 \rightsquigarrow \mathcal{C}$ and $\Sigma_0 \rightsquigarrow \Sigma$ is again an application of a basic cut-and-glue. This time, filling in defects means from the perspective

of the boundary to cut $\mathcal{C}_0$ along the $\overline{\mu}_i$ and cap off holes with discs. The operation lifts to Σ_0. Applying Lemma 14 and noting that the $\overline{\mu}_i$ have two distinct lifts (so no branch points are created/destroyed) yields the sequence (2.15), with identically vanishing homology.

For part (c), we construct H as follows. Observe that $H_1^-(\Sigma_0)/K'$ is torsion-free (because $\langle G', n\alpha\rangle = 0 \Leftrightarrow \langle G', \alpha\rangle = 0$), so we can decompose $H_1^-(\Sigma_0) = K' \oplus K'^{\perp}$ for some (non-unique!) sublattice $K'^{\perp} \simeq H_1^-(\Sigma_0)/K'$. Let $\widetilde{H} := K \oplus K'^{\perp}$, and note that $H_1^-(\Sigma_0)/\widetilde{H} \simeq K'/K \simeq G'/G$ is 2-torsion. Then we set $H := g(\widetilde{H}) \subset H_1^-(\Sigma)$. Since the cokernel of the injection g is 2-torsion, $H_1^-(\Sigma)/H$ contains at most 2- and 4-torsion. The isomorphisms (2.16) follow from $H_1^-(\Sigma) \simeq K'/G' \simeq K/G$ together with the fact that g is injective (and preserves the intersection form). $\square$

Remark. We emphasize, following Appendix A.4, that g can explicitly be defined by representing a cycle $[\gamma] \in H_1^-(\Sigma_0)$ by a curve (or curves) γ away from the gluing regions, then including γ as curves in Σ and passing to homology. Similarly, q is defined by representing $[\gamma] \in K' \subset H_1^-(\Sigma_0)$ by curve(s) γ disjoint from defects, then filling in the defects and passing to homology.

Gluing can also be extended in a straightforward manner to unit tangent bundles and twisted homology. Namely, the gluing of surfaces $\Sigma \rightsquigarrow \Sigma_0$ and $\Sigma_0 \rightsquigarrow \Sigma'$ extends uniquely to unit tangent bundles ($T_1\Sigma \rightsquigarrow T_1\Sigma_0$, $T_1\Sigma_0 \rightsquigarrow T_1\Sigma'$) by performing all cut/glue operations smoothly, *i.e.* cutting and gluing discs (etc.) along circles without introducing any kinks. Then, for example, the (even) classes $[\nu_i] - u \in \widetilde{H}_1(\Sigma_0)$ represented by non-intersecting smooth curves $\nu_i \subset \Sigma_0$ become trivial when included in $\widetilde{H}_1(\Sigma)$; and the classes $[\mu_j^{\pm}] - u \in \widetilde{H}_1(\Sigma)$ represented by the $\mu_j^{\pm}$ become trivial in $\widetilde{H}_1(\Sigma')$ once defects are filled in. Let us therefore define

$$\begin{array}{ll} \widetilde{G}_0 \subset \widetilde{H}_1(\Sigma_0) & \text{subgroup generated by } \mu_j^{\pm} - u \\ \widetilde{G}' := \ker P_+|_{\tilde{G}_0} \subset \widetilde{H}_1^-(\Sigma_0) & \\ \widetilde{G} := \operatorname{im} P_-|_{\tilde{G}_0} \subset \widetilde{H}_1^-(\Sigma_0) & \text{subgroup generated by } \mu_j := \mu_j^+ - \mu_j^- \,, \end{array} \tag{2.17}$$

where we now use 'μ_j' and '$\mu_j^{\pm}$' to denote the classes of the canonical lifts of the gluing curves to $T_1\Sigma_0$. Let $\widetilde{K}' = \ker \langle \widetilde{G}', * \rangle|_{\widetilde{H}_1^-(\Sigma_0)}$.

We may define a lift $\widetilde{g} : \widetilde{H}_1^-(\Sigma_0) \to \widetilde{H}_1^-(\Sigma)$ of the map g by representing cycles $[\gamma] \in \widetilde{H}_1^-(\Sigma_0)$ by curves $\gamma \subset T_1\Sigma_0$ away from the gluing regions, including them as curves in $T_1\Sigma$ and passing to homology. We define a lift $\widetilde{q} : K' \to \widetilde{H}_1^-(\Sigma')$ of the map q the same way, by representing cycles $[\gamma] \in K'$ by curves $\gamma \subset T_1\Sigma_0$ disjoint from the defects, filling in the defects, and passing to homology. Observe that $\widetilde{g}(u) = u$ and $\widetilde{q}(u) = u$ (*i.e.* these maps preserve the fiber class). Then

Proposition 1'

a) The map $\widetilde{g}$ is an injection of finite cokernel, and $\widetilde{q}$ is a surjection with kernel $\widetilde{G}'$. They

fit into commutative diagrams

$$\begin{array}{ccccccccc}
 & & 0 & & 0 & & & & \\
 & & \downarrow & & \downarrow & & & & \\
 & 0 & \mathbb{Z}_2 & = & \mathbb{Z}_2 & & & & \\
 & \downarrow & \downarrow & & \downarrow & & & & \\
0 & \to \widetilde{G}' & \overset{\tilde{i}}{\hookrightarrow} & \widetilde{K}' & \overset{\tilde{q}}{\to} & \widetilde{H}_1^-(\Sigma',\mathbb{Z}) & \to & 0 & \\
 & p\downarrow & & p\downarrow & & p\downarrow & & & \\
0 & \to G' & \overset{i}{\hookrightarrow} & K' & \overset{q}{\to} & H_1^-(\Sigma',\mathbb{Z}) & \to & 0\,, & \\
 & \downarrow & & \downarrow & & \downarrow & & & \\
 & 0 & & 0 & & 0 & & &
\end{array}
\qquad
\begin{array}{ccc}
0 & & 0 \\
\downarrow & & \downarrow \\
\mathbb{Z}_2 & = & \mathbb{Z}_2 \\
\downarrow & & \downarrow \\
\widetilde{H}_1^-(\Sigma_0) & \overset{\tilde{g}}{\hookrightarrow} & \widetilde{H}_1^-(\Sigma) \\
p\downarrow & & p\downarrow \\
H_1^-(\Sigma_0) & \overset{g}{\hookrightarrow} & H_1^-(\Sigma)\,. \\
\downarrow & & \downarrow \\
0 & & 0
\end{array}
\tag{2.18}$$

with all columns and all rows on the left exact.

b) Given any $\widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)$ such that $\widetilde{G} \subset \widetilde{K} \subset \widetilde{K}'$ and $\widetilde{K}/\widetilde{G} = \widetilde{K}'/\widetilde{G}'$, the quotient $\widetilde{K}'/\widetilde{K}$ is at most 2-torsion, and there exists $\widetilde{H} \subset \widetilde{H}_1^-(\Sigma)$ such that $\widetilde{H}_1^-(\Sigma)/\widetilde{H}$ is finite (at most 2- and 4-torsion), and $\widetilde{H}_1^-(\Sigma') = \widetilde{H}_1^-(\Sigma)//g(\widetilde{G}) = g(\widetilde{K})/g(\widetilde{G})$. If $\widetilde{K}$ extends a group K as in Prop 1c, then $\widetilde{H}$ can be chosen to extend H appearing there.

Proof. Part (a) is a straightforward diagram chase. On the RHS, the columns are exact by definition of twisted homology. The top square commutes because $\tilde{g}(u) = u$, and the bottom square commutes by comparing definitions of $\tilde{g}$ and g. Injectivity of g and $\tilde{g}(u) = u$ imply injectivity of $\tilde{g}$.

On the LHS, the right column is exact by definition, and the middle column is exact because $\widetilde{K}' = \langle \widetilde{G}', * \rangle|_{\tilde{H}_1^-(\Sigma_0)}$ contains u, hence is a $\mathbb{Z}_2$ extension of $K' = \langle G', * \rangle|_{H_1^-(\Sigma_0)}$. The top square commutes because $\tilde{q}(u) = u$, and $\tilde{q}$ is surjective because q and the p's are surjective. The (bottom-)right square commutes by comparing definitions of $\tilde{q}$ and q. Obviously the inclusion $\tilde{i}$ is injective, and the left square commutes by comparing definitions of $\widetilde{G}', \widetilde{K}'$ and G', K'. Exactness of the bottom row at K' implies that $\ker \tilde{q}$ is at most a $\mathbb{Z}_2$ extensions of $\operatorname{im} \tilde{i}$; but since $\tilde{q}(u) = u$ we must have $\ker \tilde{q} = \operatorname{im} \tilde{i}$. Hence the middle row is exact, and in particular $\widetilde{G}'$ does not contain the fiber class u, whence $p : \widetilde{G}' \to G'$ is an isomorphism.

Part (b) follows by repeating the proof of Prop. 1c. □

3 Paths on the small boundary

In this section we introduce an algebra $\mathbf{P}$ of paths on the boundary of a framed, triangulated 3-manifold. It is inspired by a construction of Neumann in [29] (and earlier [12]). This algebra is an extremely useful tool for describing cycles in twisted homology $\widetilde{H}_1^-(\Sigma)$, keeping track of these cycles under gluing operations, and ultimately relating the cycles to functions on the non-abelian moduli space $\mathcal{X}[\mathcal{C}]$.

3.1 Paths and generators for $H_1^-(\Sigma)$

Let $\Sigma \xrightarrow{\pi} \mathcal{C}$ be the (oriented) canonical cover of the boundary of a framed 3-manifold M (with or without defects), with a fixed 2d triangulation $\mathbf{t}_{2\mathrm{d}}$ of $\mathcal{C}_{\mathrm{big}}$. Let Σ^+ and Σ^- denote the

two sheets of Σ, constructed from a branching graph induced by $\mathbf{t}_{2d}$ as in Section 2.4 (so $\Sigma^+ \simeq \Sigma^- \simeq \mathcal{C}\backslash\Gamma_{br}$).

Let E° denote the disjoint union of the interiors all small edges of hexagons in $\mathbf{t}_{2d}$. Note that all the small edges lie on the interface $\partial\mathcal{C}_{small} \cap \partial\mathcal{C}_{big}$ between big and small parts of $\mathcal{C}$. We define $\mathbf{P}$ (depending on $\mathbf{t}_{2d}$) as the relative homology group

$$\mathbf{P} := H_1(\mathcal{C}_{small}, E^\circ)\,. \tag{3.1}$$

Thus, $\mathbf{P}$ is generated by oriented paths $\mathfrak{p}$ on $\mathcal{C}_{small}$: either closed paths, or paths whose endpoints lie in the interior of one of the small edges of $\mathbf{t}_{2d}$ (and therefore on the boundary of a distinguished hexagonal face of $\mathbf{t}_{2d}$).

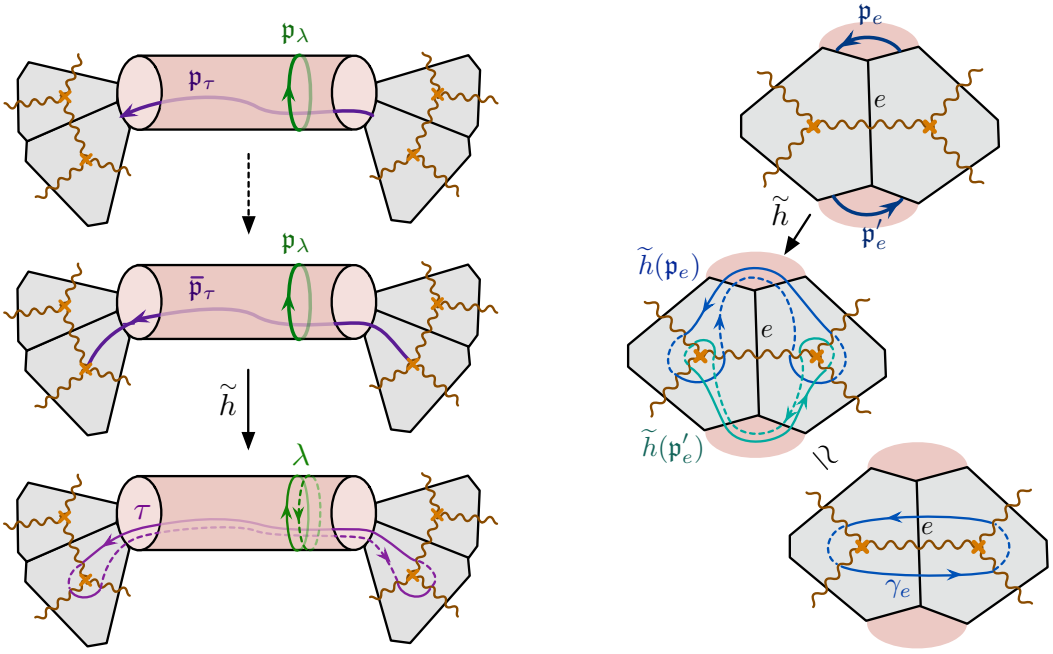

Figure 11. Examples of the map $h : \mathbf{P}[\mathbf{t}_{2d}] \to H_1^-(\Sigma)$. Left: a closed path $\mathfrak{p}(\lambda)$ surrounding a small annulus and an open path $\mathfrak{p}(\tau)$ traversing it. Right: two paths mapping equivalently to the edge cycle γ_e. For convenience in visualizing and counting intersection numbers, we slightly deform the curves representing $h(\mathfrak{p})$ for open $\mathfrak{p}$ away from branch points.

There is a homomorphism

$$\widetilde{h} \;:\; \mathbf{P} \to \widetilde{H}_1^-(\Sigma) \tag{3.2}$$

defined as follows. If $[\mathfrak{p}] \in \mathbf{P}$ can be represented by a smooth, closed curve $\mathfrak{p}$, we use the fact that there is no branching along the small boundary define lifts $\mathfrak{p}^\pm$ to $\Sigma^\pm \subset \Sigma$, and then (canonically) to $T_1\Sigma$. We set $\widetilde{h}([\mathfrak{p}]) := \ell^-(\mathfrak{p}) = [\mathfrak{p}^+] - [\mathfrak{p}^-]$ to be the class of the "odd double lift" of $\mathfrak{p}$. Note that this does not depend on which smooth representative of the curve $\mathfrak{p}$ is used, since changing the representative (as in Figure 4) can only modify $\widetilde{h}([\mathfrak{p}])$ by an even multiple of u.

If $[\mathfrak{p}] \in \mathbf{P}$ can be represented by a smooth open path $\mathfrak{p}$, then we can extend $\mathfrak{p}$ slightly to a path $\overline{\mathfrak{p}}$ that starts and ends at the branch points on the faces at the ends of $\mathfrak{p}$ (as on the left of Figure 11). We take lifts $\overline{\mathfrak{p}}^{\pm}$ to $\Sigma^{\pm}$, and connect them at the branch points to form a smooth closed curve $\mathfrak{p}^+ \circ \mathfrak{p}^-$; we then define $\widetilde{h}([\mathfrak{p}]) := \ell^-(\mathfrak{p}) = [\mathfrak{p}^+ \circ \mathfrak{p}^-]$. For convenience, we can deform the curve $\mathfrak{p}^+ \circ \mathfrak{p}^-$ so that it passes around the branch points, as on the bottom-left of Figure 11. An easy exercise verifies that it does not matter whether we pass clockwise or counter-clockwise (top of Figure 12).

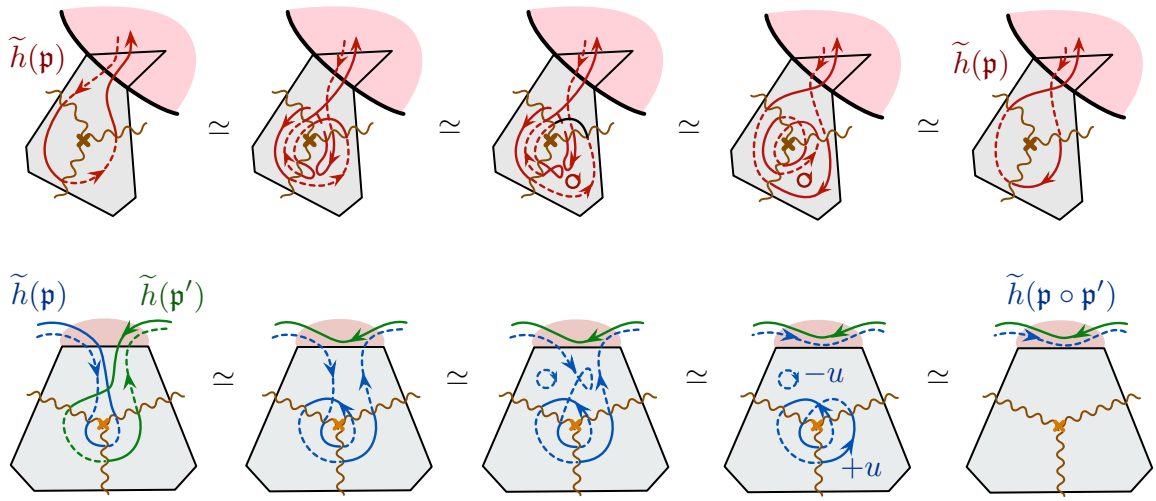

Figure 12. Top: verification that it does not matter how $\widetilde{h}(\mathfrak{p})$ is chosen to wind around branch points. Bottom: proof that $\widetilde{h}(\mathfrak{p}) + \widetilde{h}(\mathfrak{p}') = \widetilde{h}(\mathfrak{p} \circ \mathfrak{p}') = \widetilde{h}(\mathfrak{p} + \mathfrak{p}')$ in $\widetilde{H}_1^-(\Sigma)$, assuring that $\widetilde{h}$ is well defined. We represent elements of $\widetilde{H}_1(\Sigma)$ by smooth curves on Σ, keeping track of the fiber class u by means of small loops. In each case, corrections involving the fiber class cancel out at the end.

We extend $\widetilde{h}$ to all of $\mathbf{P}$ by linearity. For the map $\widetilde{h}$ to be a well-defined homomorphism, we must check that $\widetilde{h}(-\mathfrak{p}) = -\widetilde{h}(\mathfrak{p})$, which follows easily from our definitions, and that

$$\widetilde{h}(\mathfrak{p} \circ \mathfrak{p}') = \widetilde{h}(\mathfrak{p}) + \widetilde{h}(\mathfrak{p}')\,, \tag{3.3}$$

where $\mathfrak{p} \circ \mathfrak{p}'$ is a smooth path homotopic to the concatenation of $\mathfrak{p}$ and $\mathfrak{p}'$ (possibly a cyclic concatenation, in which case $\mathfrak{p} \circ \mathfrak{p}'$ is closed). The proof of (3.3) is entirely local (depending on behavior at the endpoints of $\mathfrak{p}$ and $\mathfrak{p}'$), and illustrated on the bottom of Figure 12.

The map $\widetilde{h}$ has a nontrivial kernel. If $\mathfrak{p}_e$ and $\mathfrak{p}'_e$ are two paths circling counter-clockwise around the two ends of a big edge e of $\mathbf{t}_{2\mathrm{d}}$, as on right of Figure 11, then

$$\widetilde{h}(\mathfrak{p}_e) = \widetilde{h}(\mathfrak{p}'_e)\,. \tag{3.4}$$

Let $\mathbf{P}_E = \langle \mathfrak{p}_e - \mathfrak{p}'_e \rangle_{\text{big edges } e}$ be the subgroup of $\mathbf{P}$ generated by the differences of paths associated to all edges of $\mathbf{t}_{2d}$. Clearly $\mathbf{P}_E \subset \ker \widetilde{h}$. In fact,

Lemma 3 *If M has no defects,* $\ker \widetilde{h} = \mathbf{P}_E$, *and* $\operatorname{im} \widetilde{h} \simeq H_1^-(\Sigma) \subset \widetilde{H}_1^-(\Sigma)$. *In other words,*

$$0 \to \mathbf{P}_E \hookrightarrow \mathbf{P} \xrightarrow{\tilde{h}} H_1^-(\Sigma) \to 0 \tag{3.5}$$

is exact. Thus $\widetilde{h}$ provides a splitting $\widetilde{H}_1^-(\Sigma) \simeq H_1^-(\Sigma) \oplus \mathbb{Z}_2$; and we may trivially extend $\widetilde{h}$ to a surjection $\widetilde{h} : (\mathbf{P} \oplus \mathbb{Z}_2) \twoheadrightarrow \widetilde{H}_1^-(\Sigma)$, defining $\widetilde{h}(u) = u$ for the generator u of $\mathbb{Z}_2$:

$$0 \to \mathbf{P}_E \hookrightarrow \mathbf{P} \oplus \mathbb{Z}_2 \xrightarrow{\tilde{h}} \widetilde{H}_1^-(\Sigma) \to 0\,. \tag{3.6}$$

Proof. Consider the map $h := p \circ \widetilde{h} : \mathbf{P} \to H_1^-(\Sigma)$. Momentarily we will demonstrate that h is surjective with $\ker h = \mathbf{P}_E$, by producing explicit presentations of $\mathbf{P}$ and $H_1^-(\Sigma)$. (These explicit presentations will be used throughout the rest of the paper.) Therefore $\ker \widetilde{h} \subset \mathbf{P}_E$, whence $\ker \widetilde{h} = \mathbf{P}_E$, and the rest of the Lemma follows. □

3.2 Generators of $H_1^-(\Sigma)$ and intersection form

Suppose that the triangulated framed 3-manifold M has no defects, and that $\pi_1(\mathcal{C}_{\text{small}}$ is abelian. In other words, the connected components of $\mathcal{C}_{\text{small}}$ are discs, annuli, tori, or spheres. We describe a set of generators for $H_1^-(\Sigma)$ labelled by paths $\mathfrak{p} \in \mathbf{P}$. Let $h := p \circ \widetilde{h} : \mathbf{P} \to H_1^-(\Sigma)$.

Small spheres are connected components of $\mathcal{C}$, over which both the path algebra and odd homology are trivial, so we may ignore them. Small tori are also connected components of $\mathcal{C}$. For each small torus t, choose a basis of A and B cycles, let $\mathfrak{p}_\alpha^{(t)}, \mathfrak{p}_\beta^{(t)}$ be paths representing these cycles — generating the path algebra on t — and set

$$\alpha^{(t)} = h(\mathfrak{p}_\alpha^{(t)})\,, \quad \beta^{(t)} = h(\mathfrak{p}_\beta^{(t)})\,. \tag{3.7}$$

Clearly $\alpha^{(t)}, \beta^{(t)}$ generate the part of $H_1^-(\Sigma)$ associated to the small torus.

The remaining connected components of $\mathcal{C}$ (which are all we consider now) contain big boundary $\mathcal{C}_{\text{big}}$ with holes connected to small discs or annuli. For each small annulus a let $\mathfrak{p}_\lambda^{(a)}$ be a closed path generating $H_1(a)$, and choose an open path $\mathfrak{p}_\tau^{(a)}$ traversing the annulus from one end to the other (Figure 11, left). Set

$$\lambda^{(a)} = h(\mathfrak{p}_\lambda^{(a)})\,, \quad \tau^{(a)} = h(\mathfrak{p}_\tau^{(a)})\,. \tag{3.8a}$$

For each edge e of the big-boundary triangulation $\mathfrak{t}_{2\text{d}}$, let $\mathfrak{p}_e$ and $\mathfrak{p}'_e$ be the open paths on $\mathcal{C}_{\text{small}}$ winding counter-clockwise around the two ends of e (Figure 11, right), and set

$$\gamma_e = h(\mathfrak{p}_e) = h(\mathfrak{p}'_e)\,. \tag{3.8b}$$

There are some relations among these elements. It is easy to see the path algebra $\mathbf{P}$ for small discs and annuli is generated by $\mathfrak{p}_e, \mathfrak{p}'_e, \mathfrak{p}_\lambda^{(a)}, \mathfrak{p}_\tau^{(a)}$ subject to the relations that the sum of

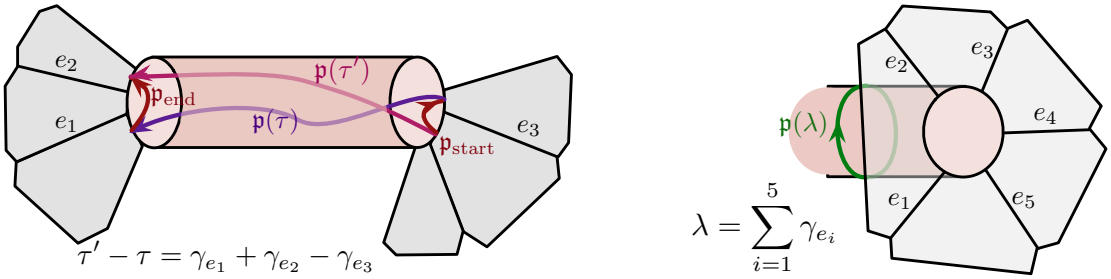

Figure 13. Examples of relations among elements $h(\mathfrak{p})$. On the left, we move the endpoints of $\mathfrak{p}_\tau$ by concatenating with (sums of) edge-paths, $\mathfrak{p}'_\tau = \mathfrak{p}_{\text{end}} \circ \mathfrak{p}_\tau \circ \mathfrak{p}_{\text{start}}$.

paths around the boundary of each small disc d vanishes, and the sums of paths around the two boundaries of each annulus a are both equal to $\mathfrak{p}_\lambda^{(a)}$,

$$\mathbf{P} = \left\langle \{\mathfrak{p}_\lambda^{(a)}, \mathfrak{p}_\tau^{(a)}\}_a,\, \{\mathfrak{p}_e, \mathfrak{p}'_e\}_e \,\middle|\, \sum_{\text{around } \partial\, d} (\mathfrak{p}_e \text{ or } \mathfrak{p}'_e) = 0,\ \sum_{\text{around } \partial_1 a} (\mathfrak{p}_e \text{ or } \mathfrak{p}'_e) = -\sum_{\text{around } \partial_2 a} (\mathfrak{p}_e \text{ or } \mathfrak{p}'_e) = \mathfrak{p}_\lambda^{(a)} \right\rangle . \tag{3.9}$$

Also note that any two choices of traversing paths $\mathfrak{p}_\tau^{(a)}$, $\mathfrak{p}_\tau^{(a)\prime}$ are related by adding or subtracting appropriate edge paths $\mathfrak{p}_e, \mathfrak{p}'_e$ (Figure 13, left).

Similarly, since the branching graph Γ_{br} (dual to big edges of $\mathbf{t}_{2\text{d}}$) together with the extended paths $\overline{\mathfrak{p}}_\tau^{(a)}$ forms the 1-skeleton of a cell decomposition of $\mathcal{C}$, lifting to the 1-skeleton of a cell decomposition of Σ, we find that (ignoring small spheres and tori)

$$H_1^-(\Sigma) = \left\langle \{\lambda^{(a)}, \tau^{(a)}\}_a,\, \{\gamma_e\}_e \,\middle|\, \sum_{e \text{ on } \partial d} \gamma_e = 0,\ \sum_{e \text{ on } \partial_1 a} \gamma_e = -\sum_{e \text{ on } \partial_2 a} \gamma_e = \lambda^{(a)} \right\rangle , \tag{3.10}$$

with each relation coming from a 2-cell. (Such cell decompositions with 0-cells at branch points are discussed in greater detail in Appendix A.3. See Lemma 11 there.) Applying h to the presentation (3.9) of $\mathbf{P}$ clearly produces the presentation (3.10) of $H_1^-(\Sigma)$, with kernel precisely $\mathbf{P}_E = \langle \mathfrak{p}_e - \mathfrak{p}'_e \rangle$.

By direct inspection, we also see that

Lemma 4 *The intersection form among generators of $H_1^-(\Sigma)$ (including small tori and spheres) is*

$$\begin{aligned}
&\langle \alpha^{(t)}, \beta^{(t')} \rangle = 2\langle \mathfrak{p}_\alpha^{(t)}, \mathfrak{p}_\beta^{(t')} \rangle = 2\,\delta_{tt'} \\
&\langle \tau^{(a)}, \lambda^{(a')} \rangle = 2\langle \mathfrak{p}_\tau^{(a)}, \mathfrak{p}_\lambda^{(a')} \rangle = \pm 2\,\delta_{aa'} \\
&\langle \tau^{(a)}, \gamma_e \rangle = \pm 1 \quad \textit{for the six } e \textit{ on faces adjacent to } \partial \mathfrak{p}_\tau^{(a)} \\
&\langle \gamma_e, \gamma_{e'} \rangle = \#\textit{ faces shared by } e, e' ,
\end{aligned} \tag{3.11}$$

where the signs in $\langle \tau, \lambda \rangle$ depends on orientation, the signs in $\langle \tau, \gamma_e \rangle$ are shown in Figure 14, and the faces in $\langle \gamma_e, \gamma'_e \rangle$ are also counted with orientation, $+1$ (-1) if e' occurs counterclockwise (clockwise) after e.

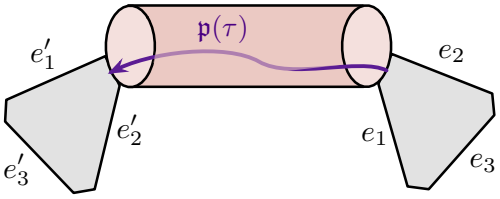

$$\langle \tau, \gamma_{e_1} \rangle = \langle \tau, \gamma_{e_2} \rangle = -\langle \tau, \gamma_{e_3} \rangle = 1$$
$$\langle \tau, \gamma_{e_1'} \rangle = \langle \tau, \gamma_{e_2'} \rangle = -\langle \tau, \gamma_{e_3'} \rangle = -1$$
$$\langle \gamma_{e_1}, \gamma_{e_2} \rangle = \langle \gamma_{e_2}, \gamma_{e_3} \rangle = \langle \gamma_{e_3}, \gamma_{e_1} \rangle = 1$$

Figure 14. Intersection product among τ's and γ_e's.

We have shown that $h : \mathbf{P} \to H_1^-(\Sigma)$ is surjective with kernel $\mathbf{P}_E$ when $\pi_1(\mathcal{C}_{\text{small}})$ is abelian. The same basic argument can easily be extended to reach the same conclusion for general $\mathcal{C}_{\text{small}}$ (as long as M has no defects). For closed components of $\mathcal{C}_{\text{small}}$ of any genus g, we choose $2g$ A and B cycles to generate $\mathbf{P}$; then applying h to these cycles (*i.e.* taking odd double lifts) directly produces generators of $H_1^-(\Sigma)$. For components of $\mathcal{C}$ containing both big and small parts, we choose multiple non-intersecting paths akin to the $\mathfrak{p}_\tau$'s, traversing $\mathcal{C}_{\text{small}}$ from boundary to boundary, so that the augmentation of Γ_{br} by these paths continues to produce a 1-skeleton of a cell decomposition of $\mathcal{C}$. Then these generalized traversing paths together with $\mathfrak{p}_e, \mathfrak{p}_e'$ generate $\mathbf{P}$ (with the only relations being that the sum of $\mathfrak{p}_e, \mathfrak{p}_e'$ paths around the boundary of any component of $\mathcal{C}_{\text{small}}$ must vanish); while applying h produces generators of $H_1^-(\Sigma)$ (with the only new relations of the form $h(\mathfrak{p}_e) = h(\mathfrak{p}_e')$). This completes the proof of Lemma 3 above.

3.2.1 Example: the tetrahedron

Let us illustrate Lemma 3 when $M = \Delta$ is a truncated tetrahedron, with its canonical triangulation $\mathbf{t}_{2\text{d}}$. The cover $\Sigma_\Delta \xrightarrow{\pi} \mathcal{C}_\Delta = \partial M_\Delta$ is branched over four points (one in the center of each big hexagonal face of the tetrahedron), so

$$\mathcal{C}_\Delta \simeq \partial\Delta \simeq S^2\,, \qquad \Sigma_\Delta \simeq T^2\,, \tag{3.12}$$

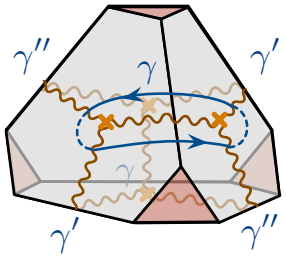

Figure 15. Six edge cycles generating $H_1^-(\Sigma_\Delta)$.

Since $\pi_* : H_1(\Sigma_\Delta) \to H_1(\mathcal{C}_\Delta)$ is trivial, all of $H_1^-(\Sigma_\Delta) = H_1(\Sigma_\Delta) = \mathbb{Z}^2$ is odd. Lemma 3 shows that $H_1^-(\Sigma_\Delta)$ is generated by six edge cycles γ_e subject to the four relations that the sum around every vertex vanishes. This implies that cycles associated to opposite edges are equal. Denoting pairs of edge cycles $\gamma, \gamma', \gamma''$ as in Figure 15, we see that

$$\begin{gathered} H_1^-(\Sigma_\Delta) = \mathbb{Z}\langle \gamma, \gamma', \gamma'' \,|\, \gamma + \gamma' + \gamma'' = 0 \rangle\,, \\ \langle \gamma, \gamma' \rangle = \langle \gamma', \gamma \rangle = \langle \gamma'', \gamma \rangle = 1\,, \end{gathered} \tag{3.13}$$

with intersection product as described by Lemma 4.

3.3 $\widetilde{G}$, $\widetilde{K}$, and the gluing map

Let's now revisit the setup of Section 2.5, where we start with a framed, triangulated 3-manifold M, glue interiors of pairs of faces to form M_0 (with defects), and fill in the defects to form M'. Let us assume that neither M nor M' have defects. We can use paths to explicitly describe subgroups $\widetilde{G} \subset \widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)$ appearing in Propositions 1, 1', and their embeddings into $\widetilde{H}_1^-(\Sigma)$ via the injection $\widetilde{g}$. Let $\mathbf{t}_{2\mathrm{d}}$ denote the induced triangulation of $\mathcal{C}_{\mathrm{big}}$ and $\mathbf{t}'_{2\mathrm{d}}$ denote the induced triangulation of $\mathcal{C}'_{\mathrm{big}} = (\mathcal{C}_0)_{\mathrm{big}}$. Let $\mathbf{P}$, $\mathbf{P}_0$, $\mathbf{P}'$ denote the algebras of paths on the respective small boundaries $\mathcal{C}_{\mathrm{small}}, (\mathcal{C}_0)_{\mathrm{small}}, \mathcal{C}'_{\mathrm{small}}$.

We begin with the subgroups of $\widetilde{H}_1^-(\Sigma_0)$. The small boundary $(\mathcal{C}_0)_{\mathrm{small}}$ of M_0 has pairs of holes at the ends of each defect I_j, *i.e.* at $(\mathcal{C}_0)_{\mathrm{small}} \cap (\mathcal{C}_0)_{\mathrm{def}}$. Let $\mathfrak{p}_{\mu_j}$ and $\mathfrak{p}'_{\mu_j}$ be paths on $(\mathcal{C}_0)_{\mathrm{small}}$ circling counter-clockwise (from the perspective of the small boundary) around the holes at the ends of I_j, as in Figure 16. Then the homology classes

$$\mu_j = \widetilde{h}(\mathfrak{p}_{\mu_j}) = \widetilde{h}(\mathfrak{p}'_{\mu_j}) \tag{3.14}$$

are precisely the cycles $\mu_j^+ - \mu_j^-$ discussed in (2.17), and generate $\widetilde{G}$. Letting $\mathbf{P}_G \subset \mathbf{P}_0$ denote the subgroup generated by $\mathfrak{p}_{\mu_j}, \mathfrak{p}'_{mu_j}$, we have

$$\widetilde{G} = \mathrm{im}\,(\widetilde{h} : \mathbf{P}_G \to \widetilde{H}_1^-(\Sigma_0))\,. \tag{3.15}$$

We then define a subgroup $\widetilde{K} \supset \widetilde{G}$ as

$$\widetilde{K} = \mathrm{im}\,(\widetilde{h} : (\mathbf{P}_0 \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma_0))\,, \tag{3.16}$$

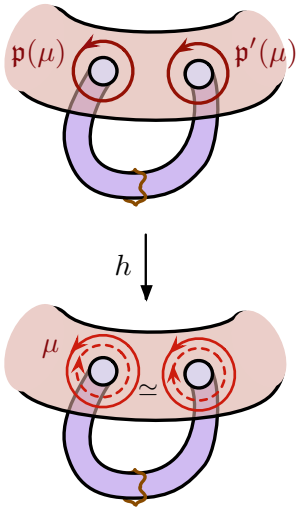

Figure 16. Generators of G

extending $\widetilde{h}$ as usual by $\mathbb{Z}_2 = \mathbb{Z}\langle u\rangle/\mathbb{Z}\langle 2u\rangle$, with $\widetilde{h}(u) = u$. Set $G = p(\widetilde{G}) = h(\mathbf{P}_G)$, $K = p(\widetilde{K}) = h(\mathbf{P}_0)$.

It is clear[7] that $\widetilde{G} \subset \widetilde{K} \subset \widetilde{K}'$, in the notation of Proposition 1'. To see that $\widetilde{K}/\widetilde{G} \simeq \widetilde{H}_1^-(\Sigma')$ so that $\widetilde{K}$ satisfies the hypotheses of part (b) of Proposition 1', we use:

Lemma 5a *The maps $\widetilde{h} : (\mathbf{P}_0 \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma_0)$ and $\widetilde{h} : (\mathbf{P}' \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma')$ fit into a commutative diagram*

$$\begin{array}{ccccccccc} 0 & \to & \mathbf{P}_G & \overset{i}{\hookrightarrow} & \mathbf{P}_0 \oplus \mathbb{Z}_2 & \overset{q_{\mathbf{P}}}{\to} & \mathbf{P}' \oplus \mathbb{Z}_2 & \to & 0 \\ & & \widetilde{h}\downarrow & & \widetilde{h}\downarrow & & \widetilde{h}\downarrow & & \\ 0 & \to & \widetilde{G} & \overset{\tilde{i}}{\hookrightarrow} & \widetilde{K} & \overset{\tilde{q}}{\to} & \widetilde{H}_1^-(\Sigma') & \to & 0\,. \end{array} \tag{3.17}$$

with top and bottom rows exact. Thus $\widetilde{K}$ satisfies the hypotheses of Prop. 1'(b).

Proof. The inclusions $\mathbf{P}_G \overset{i}{\hookrightarrow} \mathbf{P}_0$, $\widetilde{G} \overset{\tilde{i}}{\hookrightarrow} \widetilde{K}$ are obviously injective, and the first square commutes due to the definitions of $\widetilde{G}, \widetilde{K}$. On the top row, the map $q_{\mathbf{P}}$ includes a path $\mathfrak{p}$ on

[7] To see $\tilde{K} \subset \tilde{K}'$ observe that $\tilde{K} \subset \ker\langle \tilde{G}, * \rangle|_{\tilde{H}_1^-(\Sigma_0)} = \ker\langle \tilde{G}', * \rangle|_{\tilde{H}_1^-(\Sigma_0)} = \tilde{K}'$.

$(\mathcal{C}_0)_{\text{small}}$ as a path on $\mathcal{C}'_{\text{small}}$, where the holes from defects have been filled in; while $q_{\mathbf{P}}(u) := u$ for $u \in \mathbb{Z}_2$. We can invert $q_{\mathbf{P}}$ by lifting any path $\mathfrak{p}$ on $\mathcal{C}'_{\text{small}}$ to a path $\hat{\mathfrak{p}}$ on $(\mathcal{C}_0)_{\text{small}}$, choosing some way for $\hat{\mathfrak{p}}$ to pass around the defects; thus $\mathbf{P}' \simeq \mathbf{P}_0/\mathbf{P}_G$ and $\mathbf{P}' \oplus \mathbb{Z}_2 \simeq (\mathbf{P}_0 \oplus \mathbb{Z}_2)/\mathbf{P}_G$. (Alternatively, exactness of the top row follows from the long exact sequence for the pair $((\mathcal{C}_0)_{\text{small}}/E^{\circ\prime}, \mathcal{C}'_{\text{small}}/E^{\circ\prime})$ where $E^{\circ\prime}$ is the union of interiors of small edges of $\mathbf{t}'_{2d}$.) By definition of $\widetilde{K}$ and by Lemma 3 we know that all the downward maps are surjective. On the bottom row, we first observe that any cycle in $\widetilde{K}$ has zero intersection number with all defect cycles $\mu_j^{\pm} \in \widetilde{G}'$ (as in (2.17)), so we may define $\widetilde{q}$ to act by including these cycles in $\widetilde{H}_1^-(\Sigma)$. In other words, $\widetilde{q}$ is the restriction of the map in (2.18) to $\widetilde{K} \subset \widetilde{K}'$. With this definition, the second square also commutes (see Figure 17), and surjectivity of $\widetilde{q}$ on the bottom row follows from surjectivity of $q_{\mathbf{P}}$ on the top. Thus $\widetilde{K}/\widetilde{G} \simeq \widetilde{H}_1^-(\Sigma')$. $\square$

Remark. By applying p to (3.17) and killing all the $\mathbb{Z}_2$'s and fiber classes, we see that $G = p(\widetilde{G}) = h(\mathbf{P}_G)$ and $K = p(\widetilde{K}) = h(\mathbf{P}_0)$ satisfy the hypotheses of Proposition 1(c).

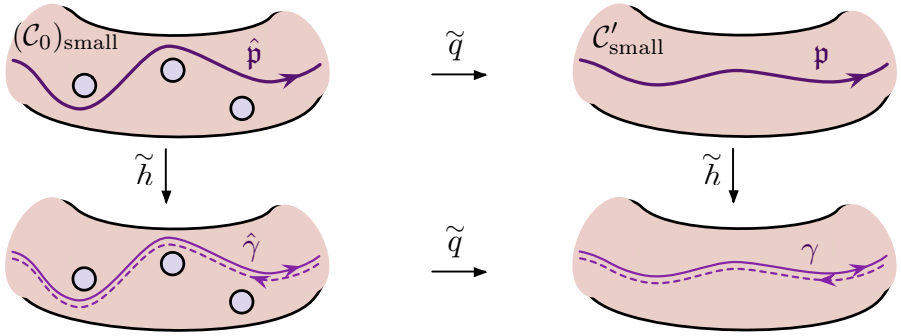

Figure 17. Commutativity of the second square in (3.17).

Explicitly, the sequences (3.17) imply that to generate K (or $\widetilde{K} = K \oplus \mathbb{Z}_2$) we may start with generators $\gamma_e, \tau, \ldots$ of $H_1^-(\Sigma')$ labelled by paths $\mathfrak{p}_e, \mathfrak{p}_\tau, \ldots$, use $q_{\mathbf{P}}^{-1}$ to lift the paths to any $\hat{\mathfrak{p}}_e, \hat{\mathfrak{p}}_\tau, \ldots \in \mathbf{P}_0$ (unique up to addition of $\mathfrak{p}_{\mu_j}$ and $\mathfrak{p}'_{\mu_j}$), adjoin the $\mathfrak{p}_{\mu_j}$'s, and then apply h to obtain generators $\hat{\gamma}_e, \hat{\tau}, \ldots$ and μ_j of K. The lifts $\hat{\gamma}_e, \hat{\tau}, \ldots$ satisfy $q(\hat{\gamma}_e) = \gamma_e$, $q(\hat{\tau}) = \tau$, etc., and are unique up to addition of μ_j's.

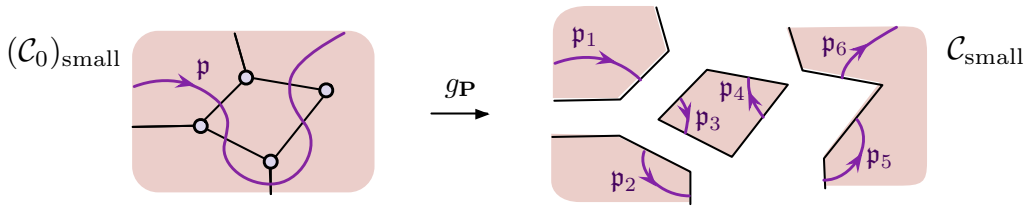

Figure 18. The cutting map acting on paths.

We may similarly use paths to find an algorithm for the action of the injection $\widetilde{g} : \widetilde{H}_1^-(\Sigma_0) \to \widetilde{H}_1^-(\Sigma)$ on $\widetilde{K}$.

Recall that $\mathbf{P}$ is a relative homology group $H_1(\mathcal{C}_{\text{small}}, E^\circ)$ where E° denotes the disjoint union of interiors of small edges of $\mathbf{t}_{2\text{d}}$. Similarly, $\mathbf{P}_0 = H_1((\mathcal{C}_0)_{\text{small}}, E^{\circ\prime})$, where $E^{\circ\prime}$ is the union of interiors of small edges of $\mathbf{t}'_{2\text{d}}$. The small edges E° have an image $\hat{E}^\circ \supset E^{\circ\prime}$ in $(\mathcal{C}_0)_{\text{small}}$ — along which $(\mathcal{C}_0)_{\text{small}}$ is cut to obtain $\mathcal{C}_{\text{small}}$. Since $E^{\circ\prime} \subset \hat{E}^\circ$, there exists a natural "cutting map" $g_{\mathbf{P}} : H_1((\mathcal{C}_0)_{\text{small}}, E^{\circ\prime}) \to H_1((\mathcal{C}_0)_{\text{small}}, \hat{E}^\circ) \simeq H_1(\mathcal{C}_{\text{small}}, E^\circ)$, or

$$g_{\mathbf{P}} : \begin{array}{ccc} \mathbf{P}_0 & \to & \mathbf{P} \\ \mathfrak{p} & \mapsto & \sum_i \mathfrak{p}_i \,. \end{array} \tag{3.18}$$

Explicitly, $g_{\mathbf{P}}$ takes a path on $(\mathcal{C}_0)_{\text{small}}$ to the sum of paths $\sqcup_i \mathfrak{p}_i = \mathfrak{p} \cap \mathcal{C}_{\text{small}}$ in its intersection with patches of $\mathcal{C}_{\text{small}}$, as in Figure 18. We also define the homomorphism

$$\text{cut} : \begin{array}{ccc} \mathbf{P}_0 & \to & \mathbb{Z}_2 = \{0, 1\} \\ \mathfrak{p} & \mapsto & \#(\mathfrak{p} \cap (E^\circ \backslash E^{\circ\prime})) \end{array} \tag{3.19}$$

to be the unoriented intersection number between a path (or sum of paths) $\mathfrak{p}$ and the small edges of $E^{\circ\prime}$, included as a subset of $(\mathcal{C}_0)_{\text{small}}$. In other words, $\text{cut}(\mathfrak{p})$ is the number of times a path $\mathfrak{p}$ is cut when the map $g_{\mathbf{P}}$ is applied. It is well defined modulo 2. (As an example, in Figure 18, $\text{cut}(\mathfrak{p}) = 5 \equiv 1$.) Then we construct

$$\widetilde{g}_{\mathbf{P}} : \begin{array}{ccc} \mathbf{P}_0 \oplus \mathbb{Z}_2 & \to & \mathbf{P} \oplus \mathbb{Z}_2 \\ \mathfrak{p} + nu & \mapsto & g_{\mathbf{P}}(\mathfrak{p}) + (\text{cut}(\mathfrak{p}) + n)u \,. \end{array} \tag{3.20}$$

Note that $\widetilde{g}_{\mathbf{P}}$ acts trivially on $u \in \mathbb{Z}_2$, but sends a path $\mathfrak{p} \in \mathbf{P}_0$ to $g_{\mathbf{P}}(\mathfrak{p}) + \text{cut}(\mathfrak{p})u \in \mathbf{P} \oplus \mathbb{Z}_2$.

Lemma 5b *The maps $\widetilde{h} : \mathbf{P}_0 \to H_1^-(\Sigma_0)$ and $\widetilde{h} : \mathbf{P} \to H_1^-(\Sigma)$ and the injection $\widetilde{g}$ fit into a commuting square*

$$\begin{array}{ccc} \mathbf{P}_0 \oplus \mathbb{Z}_2 & \overset{\widetilde{g}_{\mathbf{P}}}{\to} & \mathbf{P} \oplus \mathbb{Z}_2 \\ \widetilde{h} \downarrow & & \widetilde{h} \downarrow \\ \widetilde{H}_1^-(\Sigma_0) & \overset{g}{\hookrightarrow} & \widetilde{H}_1^-(\Sigma) \,, \end{array} \tag{3.21}$$

so that $\widetilde{g}\widetilde{h} = \widetilde{h}\widetilde{g}_{\mathbf{P}}$.

Proof. The twisted map $\widetilde{g}_{\mathbf{P}}$ was constructed precisely to ensure commutativity. The proof follows from a local argument, illustrated in Figure 19. □

Remark. Lemma 5b lets us algorithmically calculate $\widetilde{g}(\gamma) \in \widetilde{H}_1^-(\Sigma)$ for any $\gamma \in \widetilde{K} = \text{im}\left[\widetilde{h} : (\mathbf{P}_0 \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma_0)\right]$. By applying p, killing all $\mathbb{Z}_2$'s and fiber classes, we may also use the bare version of the cutting map $g_{\mathbf{P}}$ to calculate $g(\gamma) \in H_1^-(\Sigma)$ for any $\gamma \in K = \text{im}\left[h : \mathbf{P}_0 \to H_1^-(\Sigma_0)\right]$. In un-twisted homology, $gh = hg_{\mathbf{P}}$.

3.4 Detour: Neumann's chain complex

We may use the formalism developed so far to provide a new topological interpretation for the chain complex that Neumann introduced in [29] to understand the symplectic properties of the hyperbolic gluing equations. (The actual connection with gluing equations will come in the next section.)

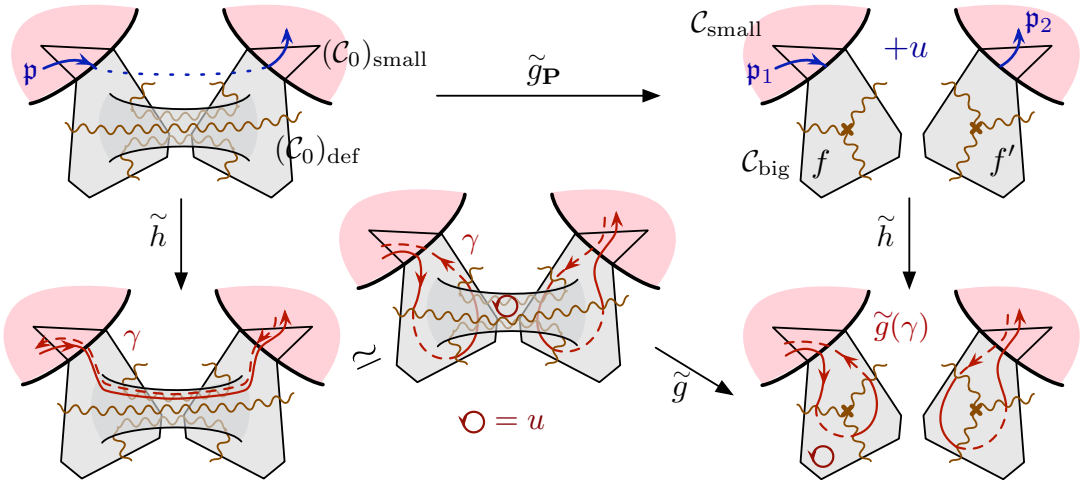

Figure 19. Commutative diagram demonstrating $\widetilde{g} \circ \widetilde{h}(\mathfrak{p}) = \widetilde{h} \circ \widetilde{g}_{\mathbf{P}}(\mathfrak{p})$. On the bottom row, the curve $\gamma = \widetilde{h}(\mathfrak{p})$ in $T_1\Sigma_0$, which crosses the gluing region, is homologous to a sum of two smooth curves on either side of the gluing region along with a curve wrapping the fiber u, represented as a small loop. This extra "twisted" correction to $\widetilde{g}(\gamma)$ is accounted for by the contribution u to $\widetilde{g}_{\mathbf{P}}(\mathfrak{p}) = g_{\mathbf{P}}(\mathfrak{p}) + \text{cut}(\mathfrak{p})u$ in the top row (here $\text{cut}(\mathfrak{p}) = 1$).

Suppose that M' is a framed 3-manifold with no defects and no big boundary (the topology of $\mathcal{C}' = \mathcal{C}'_{\text{small}}$ is unconstrained), which is glued from a collection $M = \sqcup_{i=1}^N \Delta_i$ of truncated tetrahedra. As usual we separate the gluing into two steps $M \rightsquigarrow M_0 \rightsquigarrow M'$ where M_0 has defects I_j along the edges of M'. In this section we work with untwisted homology, and untwisted maps, always projecting out the fiber class u. We build a chain complex in several steps.

First, let the $\mu_j^{\pm} \subset \Sigma_0$ be the actual curves (as in Figure 10 of Section 2.5) used to represent the generators $\mu_j = \mu_j^+ - \mu_j^-$ of $G \subset H_1^-(\Sigma_0)$, and let $N_\mu \subset \Sigma_0$ be a disjoint union of neighborhoods of these curves. (Alternatively N_μ is the lift of the defect boundary $(\mathcal{C}_0)_{\text{def}}$ to Σ.) Thus N_μ is a disjoint union of $2N$ annuli; and $H_1^-(N_\mu) \simeq \mathbb{Z}^N$ is generated by the μ_j as curves in N_μ. The image of the inclusion

$$H_1^-(N_\mu) \xrightarrow{i} H_1^-(\Sigma_0) \tag{3.22}$$

is precisely G. We may compose this with the injection $g : H_1^-(\Sigma_0) \hookrightarrow H_1^-(\Sigma) = \oplus_{i=1}^N H_1^-(\Sigma_{\Delta i})$ (of finite cokernel) to get

$$H_1^-(N_\mu) \xrightarrow{g \circ i} H_1^-(\Sigma)\,, \qquad \text{im}\,(g \circ i) = g(G)\,. \tag{3.23}$$

Next, we dualize the map $g \circ i$ with respect to the canonical *symmetric* pairing $(\mu_i, \mu_j) = \delta_{ij}$ on $H_1^-(N_\mu)$ and the intersection product $\langle\ ,\ \rangle$ on $H_1^-(\Sigma)$ to get a complex

$$H_1^-(N_\mu) \xrightarrow{g \circ i} H_1^-(\Sigma) \xrightarrow{(g \circ i)^*} H_1^-(N_\mu)\,. \tag{3.24}$$

By definition, the kernel of $(g \circ i)^*$ is $K'' := \ker \langle g(G), * \rangle|_{H_1^-(\Sigma)}$. This includes $g(K')$, where $K' = \ker \langle G, * \rangle|_{H_1^-(\Sigma_0)}$; indeed, by Proposition 1a $K''/g(K')$ is torsion. Moreover, setting $K = \mathrm{im}\,[h : \mathbf{P}_0 \to H_1^-(\Sigma_0)]$ as in Section 3.3, we obtain from Lemma 3 that $g(K) \subset g(K') \subset K''$ and $K''/g(K)$ is torsion. Thus the homology of (3.24) equals $g(K)/g(G) \simeq H_1^-(\Sigma')$ modulo torsion.

Note that $H_1^-(\Sigma')$ is isomorphic to the standard homology $H_1(\mathcal{C}') = H_1(\mathcal{C}'_{\mathrm{small}})$ since $\Sigma' \to \mathcal{C}'$ is a trivial double cover (*cf.* Lemma 12 of Appendix A.3). The isomorphism $\ell^- : H_1(\mathcal{C}') \xrightarrow{\sim} H_1^-(\Sigma')$ acts as multiplication by two on the intersection form.

We may further resolve the complex on the left using an odd version of the long exact sequence in relative homology for the pair (Σ_0, N_μ),

$$\begin{array}{ccccccccc} H_2^-(\Sigma_0) & \xrightarrow{j} & H_2^-(\Sigma_0, N_\mu) & \xrightarrow{\delta} & H_1^-(N_\mu) & \xrightarrow{g\circ i} & H_1^-(\Sigma) & \xrightarrow{(g\circ i)^*} & H_1^-(N_\mu) \\ \| & & \| & & & & & & \\ 0 & & H_2^-(\Sigma')\,. & & & & & & \end{array} \tag{3.25}$$

Note that $H_2^-(\Sigma_0) = 0$ because Σ_0 is connected; and $H_2^-(\Sigma_0, N_\mu) \simeq H_2^-(\Sigma') \simeq \mathbb{Z}^{\#\mathcal{C}'}$ has one generator for every component of $\mathcal{C}' = \partial M'$ (the odd-double-lift of the fundamental class of that component to Σ'). By Lemma 9 (Appendix A.2), the odd version of the long exact is exact modulo 2-torsion. Thus the homology of (3.25) is zero, except at $H_1^-(\Sigma)$, where it equals $H_1^-(\Sigma')$ (modulo torsion).

The complex (3.25) is identical to the one presented in Theorem 4.1 of Neumann [29], but now the various groups involved have topological meaning. The maps δ and $g \circ i$ are what Neumann calls α and β. The isomorphism $\ell^- : H_1(\mathcal{C}') \xrightarrow{\sim} H_1^-(\Sigma')$ is called δ. We could also extend the complex further to the right, as in [29], by dualizing the long exact sequence. By computing the homology of (3.25) modulo torsion, we have re-proven [29, Thm. 4.1]. Neumann goes further and carefully identifies the torsion groups as well. It would be interesting to investigate the torsion further from the perspective of odd homology.

4 Framed flat connections on boundaries

Having described gluing in terms of odd homology — both abstractly as in Proposition 1 and concretely in terms of path algebras $\mathbf{P}$ — we now proceed to describe how gluing acts on spaces of framed flat $PGL(2) := PGL(2, \mathbb{C})$ connections.

We begin by reviewing the construction of cluster-like $\mathbb{C}^*$ coordinates for moduli spaces of framed flat $PLG(2)$ connections on boundaries of framed, triangulated 3-manifolds, following [31, 44] and [32, Appendix A].[8] We show directly that the coordinates are labelled by cycles in odd homology, with Poisson, symplectic, and K_2 structures induced by the intersection form in odd homology (Propositions 2–3).

[8] Some additional subtleties and restrictions arise, especially as compared to [31] and Appendix A of [32], because we focus only on boundaries of framed 3-manifolds rather than their interiors.

We then explain gluing of framed $PGL(2)$ connections works, and derive the standard combinatorial gluing equations that appeared in [31, 32], a generalization of Thurston's gluing equations for hyperbolic 3-manifolds. Observing that these equations take the promised form (1.12), we deduce that gluing is K_2 symplectic reduction (Theorem 1).

4.1 Moduli spaces

Let $\mathcal{C} = \partial M$ be the boundary of a framed 3-manifold, as defined in Section 2.3. In particular, $\mathcal{C} = \mathcal{C}_{\text{big}} \cup_{\sqcup_i S^1_i} \mathcal{C}_{\text{small}} \cup_{\sqcup_j S^1_j} \mathcal{C}_{\text{def}}$ has a splitting into big, small, and defect parts. Let $\mathbf{t}_{2\text{d}}$ be a triangulation of the big boundary — *i.e.* a tiling of $\mathcal{C}_{\text{big}}$ by hexagons. Let $E \to \mathcal{C}$ be a trivial $\mathbb{C}^2$ bundle and $\mathbb{P}E$ its projectivization (a $\mathbb{CP}^1$ bundle). Let $\mathcal{C}^*$ be $\mathcal{C}$ with a point removed in the center of every small disc of $\mathcal{C}_{\text{small}}$ — thus $\mathcal{C}^*$ is a "punctured" boundary. We define

$$\begin{aligned}
\mathcal{X}[\mathcal{C}] :=& \{\text{framed flat } PGL(2) \text{ connections on } E \to \mathcal{C}^* \\
&\qquad \text{with unipotent monodromy at punctures}\} \\
:=& \{\text{flat } PGL(2) \text{ connections on } E \to \mathcal{C}^* \\
&\qquad \text{with unipotent monodromy at punctures} \\
&\qquad \text{and a choice of invariant flat line on each component of } \mathcal{C}_{\text{small}}\} \\
=& \{\text{flat connections on } \mathbb{P}E \to \mathcal{C}^* \\
&\qquad \text{with unipotent monodromy at punctures} \\
&\qquad \text{and a choice of global flat section on each component of } \mathcal{C}_{\text{small}}\}\,.
\end{aligned} \tag{4.1}$$

We emphasize that these moduli spaces are defined modulo $PGL(2)$ gauge isomorphism. As stated, the framing data is a choice of invariant flat line in the $\mathbb{C}^2$ bundle along the small boundary, or a flat section of the projectivized bundle.

If M has no defects, and $\pi_1(\mathcal{C}_{\text{small}})$ is abelian (*i.e.* $\mathcal{C}_{\text{small}}$ contains only small discs, annuli, tori, and spheres), then $\mathcal{X}[\mathcal{C}]$ has some additional properties. In this case a choice of framing always exists and is usually unique.[9] Forgetting the framing, one gets a map to the standard character variety

$$\mathcal{X}[\mathcal{C}] \overset{\text{forget}}{\to} \text{Hom}^{\text{un}}\big(\pi_1(\mathcal{C}^*),\, PGL(2)\big)/PGL(2) \tag{4.2}$$

(where "un" means we restrict puncture holonomies to be unipotent). For particular flat connections, there may be an enhanced choice of framing, which resolves some singularities in the character variety. If $\mathcal{C}$ is connected, the complex dimension at smooth points is

$$d(\mathcal{X}[\mathcal{C}]) = \begin{cases} 0 & \mathcal{C} \text{ is a small sphere} \\ 2 & \mathcal{C} \text{ is a small torus} \\ (2g-2)\dim(G) + n(\dim(G) - \text{rank}(G)) & \\ \qquad = 6g - 6 + 2n & \chi(\mathcal{C}^*) < 0\,, \end{cases} \tag{4.3}$$

where n denotes the number of small discs in $\mathcal{C}_{\text{small}}$, or the number of punctures of $\mathcal{C}^*$. (If $\mathcal{C} = \sqcup_i \mathcal{C}_i$ contains multiple connected components, the space $\mathcal{X}[\mathcal{C}] = \prod_i \mathcal{X}[\mathcal{C}_i]$ factorizes.)

[9] Conversely, if $\pi_1(\mathcal{C}_{\text{small}})$ is nonabelian, the existence of framing is highly restrictive!

At its smooth points, the standard character variety has a non-degenerate complex symplectic form given by the Atiyah-Bott formula[10]

$$\omega = \int_{\mathbb{C}^*} \mathrm{Tr}(\delta\mathcal{A} \wedge \delta\mathcal{A})\,, \tag{4.4}$$

which pulls back to a symplectic form on an open subset of $\mathcal{X}[\mathcal{C}]$ (which we also call ω). In some cases it is known how to extend ω to a symplectic form on all of $\mathcal{X}[\mathcal{C}]$. For example, if $\mathcal{C}$ consists only of big boundary and small discs (no annuli or tori), $\mathcal{X}[\mathcal{C}]$ coincides with the space studied by Fock and Goncharov [44] using cluster coordinates, and the appropriate extension of ω was defined therein.

We will not try to extend ω to a symplectic form on all of $\mathcal{X}[\mathcal{C}]$ here, though it does seem possible to do so. Instead we will work on open patches $\mathcal{P}[\mathcal{C}; \mathbf{t}_{2\mathrm{d}}] \subset \mathcal{X}[\mathcal{C}]$ of the form $(\mathbb{C}^*)^{d(\mathcal{X}[\mathcal{C}])}$, labelled by triangulations $\mathbf{t}_{2\mathrm{d}}$ (as in [44]), on which it is fairly straightforward to calculate and extend ω.

4.2 Coordinate functions

Fix a 2d triangulation $\mathbf{t}_{2\mathrm{d}}$ of $\mathcal{C}_{\mathrm{big}}$ as above. Let $\mathbf{P} = \mathbf{P}[\mathbf{t}_{2\mathrm{d}}]$ be the abelian group of paths on $\mathcal{C}_{\mathrm{small}}$ defined in Section 3. For each path $\mathfrak{p} \in \mathbf{P}$ we can define a $\mathbb{C}^*$-valued function $x_{\mathfrak{p}}$ on $\mathcal{X}[\mathcal{C}]$, which turns out to depend only on the homology class $\gamma = \widetilde{h}(\mathfrak{p}) \in H_1^-(\Sigma) \subset \widetilde{H}_1^-(\Sigma)$. Throughout this section, we use the splitting induced by $\widetilde{h}$ (Lemma 3) to view $H_1^-(\Sigma)$ as a subgroup of $\widetilde{H}_1^-(\Sigma)$. We proceed as follows.

For every $\mathfrak{p} \in \mathbf{P}$ represented by a closed path and $\mathcal{A} \in \mathcal{X}[\mathcal{C}]$, set

$$x_{\mathfrak{p}} := \text{squared-eigenvalue of } \mathrm{Hol}_{\mathfrak{p}}(\mathcal{A}) \text{ corresponding to the framing eigenline}\,. \tag{4.5}$$

This makes sense since $\mathfrak{p}$ lies entirely on a single small boundary component, and the framing of $\mathcal{A}$ specifies a distinguished eigenline on that component. For example, on a small torus t the A- and B-cycle paths $\mathfrak{p}_\alpha^{(t)}, \mathfrak{p}_\beta^{(t)}$ give rise to

$$x_\alpha^{(t)} := x_{\mathfrak{p}_\alpha^{(t)}}\,, \qquad x_\beta^{(t)} := x_{\mathfrak{p}_\beta^{(t)}}\,, \tag{4.6}$$

which are the squares of holonomy eigenvalues along respective cycles. On every small annulus a there is a "length" function

$$x_\lambda^{(a)} := x_{\mathfrak{p}_\lambda^{(a)}} \tag{4.7}$$

measuring the holonomy eigenvalue around the girth of the annulus (Figure 20, left). For a small disc, the path $\mathfrak{p}$ around the circumference of the disc gives $x_{\mathfrak{p}} = 1$ since the the holonomy must be unipotent there; this is consistent with the fact that $\mathfrak{p}$ is contractible.

If $\mathfrak{p}$ is represented by an open path, it starts and ends on the boundaries of some big hexagons f, f', which are faces of $\mathbf{t}_{2\mathrm{d}}$. For a given framed flat connection $\mathcal{A} \in \mathcal{X}[\mathcal{C}]$, let k

[10] Technically, (4.4) is a formula on the space of all connections; the symplectic form on the character variety obtained by a symplectic reduction of (4.4) as in [39].

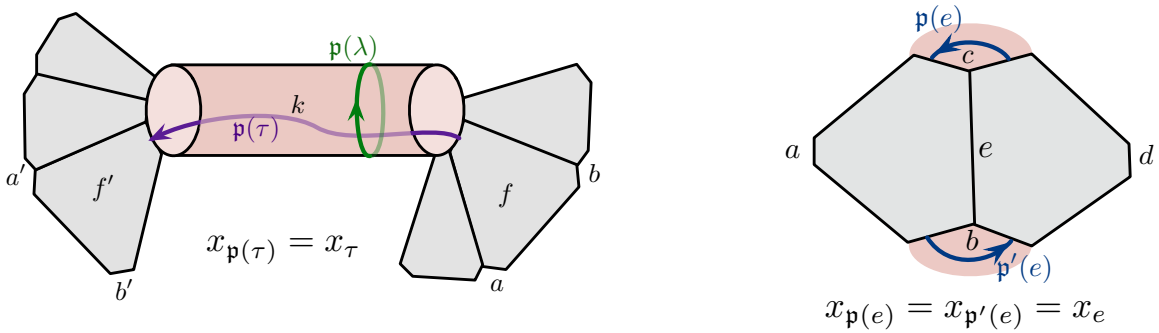

Figure 20. Left: configuration of lines that defines the function $x_{\mathfrak{p}}$ for an open path (here for the "twist" of an annulus). Right: specialization of this configuration for paths that surround a big-boundary edge.

denote the framing line of $\mathcal{A}$ along $\mathfrak{p}$. Let a, b (respectively, a', b') denote the framing lines at the two small edges of f (respectively, f') disjoint from $\mathfrak{p}$, oriented as at the ends of $\mathfrak{p}(\tau)$ in Figure 20. We can unambiguously parallel-transport a, b, a', b' and k to a single fiber of $E \to \mathcal{C}$ above (say) the midpoint of $\mathfrak{p}$. For example, we first transport a, b into the hexagon f, then transport them to the start of $\mathfrak{p}$, then along $\mathfrak{p}$; and similarly for a', b'. Working in a single fiber, we define the function $x_{\mathfrak{p}}$ to be the generalized cross-ratio

$$x_{\mathfrak{p}} = \frac{\langle a \wedge k\rangle\langle b \wedge k\rangle}{\langle a \wedge b\rangle} \cdot \frac{\langle a' \wedge b'\rangle}{\langle a' \wedge k\rangle\langle b' \wedge k\rangle} \in \mathbb{C}^* . \tag{4.8}$$

Here $\langle * \wedge * \rangle$ is a skew-symmetric volume form in the fiber of E. To calculate (4.8) one chooses any vectors in the lines a, b, a', b', k and applies the volume form. The result is independent of the normalization of these vectors, independent of the volume form, and invariant under a $PGL(2)$ action on the fiber — thus independent of the precise point along $\mathfrak{p}$ at which we compare lines. (Equivalently, working in a projectivized $\mathbb{CP}^1$ bundle, one can think of (4.8) as a generalized cross-ratio of five points in $\mathbb{CP}^1$: $(a-b)(a'-k)(b'-k)/[(a-k)(b-k)(a'-b')]$.) Note that the cross-ratio (4.8) only makes sense at points of $\mathcal{X}[\mathcal{C}]$ where all pairs of lines a, b, a', b', k being compared are independent — *i.e.* the configuration of lines is in "general position."

Definition 3 *Let $\mathcal{P}[\mathcal{C}; \mathbf{t}_{2d}] \subset \mathcal{X}[\mathcal{C}]$ be the algebraically open subset on which functions $x_{\mathfrak{p}} \in \mathbb{C}^*$ are well defined for all $\mathfrak{p} \in \mathbf{P}[\mathbf{t}_{2d}]$ represented by open paths (meaning all pairs of lines being compared to form $x_{\mathfrak{p}}$ are in general position), with the additional restriction[11] that for each connected component of $\mathcal{C}_{\text{small}}$ the holonomies around non-contractible cycles of*

[11] This is largely a technical restriction, necessary for non-degeneracy to hold in Prop. 2d below. The issue is that unipotency does not uniquely specify the conjugacy class of a holonomy matrix — it might be either trivial or parabolic. From a 3d perspective, this conjugacy class is fixed (and generically parabolic), so restricting to non-unipotent holonomy was not necessary in (*e.g.*) Appendix A.3 of [32], analogous constructions in [31], or many years of hyperbolic geometry.

that component are not simultaneously unipotent. (For example, $x_\lambda^{(a)} \neq 1$ for annuli and $(x_\alpha^{(t)}, x_\beta^{(t)}) \neq (1,1)$ for tori.) Often we simply write $\mathcal{P}[\mathcal{C}]$.

Remark. For components of $\mathcal{C}$ consisting of big boundary with holes filled by discs, $\mathcal{P}[\mathcal{C}; \mathbf{t}_{2d}]$ coincides with a cluster-coordinate chart of [44].

Proposition 2

a) For both open and closed paths $x_{-\mathfrak{p}} = x_{\mathfrak{p}}^{-1}$, and under concatenation $x_{\mathfrak{p}\circ\mathfrak{p}'} = x_{\mathfrak{p}} x_{\mathfrak{p}'}$. Therefore, we may extend path functions to all $\mathfrak{p} \in \mathbf{P}$ by linearity, defining $x_{\mathfrak{p}+\mathfrak{p}'} := x_{\mathfrak{p}} x_{\mathfrak{p}'}$, to obtain a map

$$x : \begin{array}{ccc} \mathcal{P}[\mathcal{C}] \times \mathbf{P} & \to & \mathbb{C}^* \\ \mathfrak{p} & \mapsto & x_{\mathfrak{p}} \,, \end{array} \tag{4.9}$$

that's a homomorphism on the second factor.

b) The function $x_{\mathfrak{p}}$ depends only on the homology class $\gamma = \widetilde{h}(\mathfrak{p}) \in H_1^-(\Sigma) \subset \widetilde{H}_1^-(\Sigma)$.

c) If M has no defects, the surjection $\widetilde{h} : \mathbf{P} \twoheadrightarrow H_1^-(\Sigma)$ from (3.2) provides a map

$$x : \begin{array}{ccc} \mathcal{P}[\mathcal{C}] \times H_1^-(\Sigma) & \to & \mathbb{C}^* \\ \gamma & \mapsto & x_\gamma \,, \end{array} \tag{4.10}$$

that's a homomorphism on the second factor. By defining $x_u := -1$, we may extend (4.10) to a map $x : \mathcal{P}[\mathcal{C}] \times \widetilde{H}_1^-(\Sigma) \to \mathbb{C}^*$.

d) If in addition $\pi_1(\mathcal{C}_{\text{small}})$ is abelian[12]*, the map (4.10) is non-degenerate in the sense that any basis $\{\gamma_i\}_{i=1}^d$ for $H_1^-(\Sigma) \subset \widetilde{H}_1^-(\Sigma)$ provides an injection*

$$(x_{\gamma_i}, ..., x_{\gamma_d}) : \mathcal{P}[\mathcal{C}] \hookrightarrow (\mathbb{C}^*)^{\operatorname{rank} H_1^-(\Sigma)} \,, \tag{4.11}$$

with image the subset where $x_\lambda^{(a)} \neq 1$, $(x_\alpha^{(t)}, x_\beta^{(t)}) \neq (1,1)$.

e) With the assumptions in (d), $\operatorname{rank} H_1^-(\Sigma) = d(\mathcal{X}[\mathcal{C}])$, so $\mathcal{P}[\mathcal{C}] \subset \mathcal{X}[\mathcal{C}]$ is a subset of maximal dimension, with coordinates $\{x_\gamma\}_{\gamma \in H_1^-(\Sigma)}$, isomorphic to $(\mathbb{C}^*)^{d(\mathcal{X}[\mathcal{C}])}\big|_{x_\lambda^{(a)} \neq 1,\, (x_\alpha^{(t)}, x_\beta^{(t)}) \neq (1,1)}$.

Proof. For part (a), $x_{-\mathfrak{p}} = x_{\mathfrak{p}}^{-1}$ is clear from the definitions and $x_{\mathfrak{p}\circ\mathfrak{p}'} = x_{\mathfrak{p}} x_{\mathfrak{p}'}$ follows from from a straightforward local calculation. When concatenating $\mathfrak{p} \circ \mathfrak{p}'$ to produce another open path, numerators and denominators of (4.8) cancel out to ensure $x_{\mathfrak{p}\circ\mathfrak{p}'} = x_{\mathfrak{p}} x_{\mathfrak{p}'}$. When producing a closed path, all numerators and denominators cancel up to an overall factor, which is precisely the square of the holonomy eigenvalue along the closed path.

For part (b), observe that of $\mathfrak{p}_e, \mathfrak{p}'_e$ are paths around the endpoint of a big edge e of $\mathbf{t}_{2d}$ then $x_{\mathfrak{p}_e} = x_{\mathfrak{p}'_e}$ (see (4.13) below). If M has no defects, the claim follows from Lemma 3. If M has defects, we must do a bit more work. Let M' be the framed 3-manifold obtained by filling in the defects of M, and consider the composition $q \circ h : \mathbf{P} \to H_1^-(\Sigma')$ as in (3.17)

[12] This assumption may be lifted, by using a more general version of the reconstruction procedure (with unipotent modifications to remove extraneous punctures) from Appendix B.

(we have trivially factored out a $\mathbb{Z}_2$, using untwisted versions of these maps). The kernel of this map is $\mathbf{P}_G \oplus \mathbf{P}_E$, where $\mathbf{P}_E = \langle \mathfrak{p}_e - \mathfrak{p}'_e \rangle$ are the relations of Lemma 3. Thus the kernel $\mathbf{P}_{\rm def}$ of h is a subgroup $\mathbf{P}_E \subset \mathbf{P}_{\rm def} \subset \mathbf{P}_G \oplus \mathbf{P}_E$. Indeed, $\mathbf{P}_{\rm def}$ is generated by elements $\langle \mathfrak{p}_{\mu_j} - \mathfrak{p}'_{\mu_j} \rangle$ for all defects I_j, corresponding to $h(\mathfrak{p}_{\mu_j}) - h(\mathfrak{p}'_{\mu_j}) = 0$ in (3.14), in addition to the edge relations in $\mathbf{P}_E$. Thus, for the claim of part (b), we must show that $x_{\mathfrak{p}_{\mu_j}} = x_{\mathfrak{p}'_{\mu_j}}$ for all defects I_j. This holds because the closed paths $\mathfrak{p}_{\mu_j}$ and $\mathfrak{p}'_{\mu_j}$ are homotopic on $\mathcal{C}$, so the eigenvalues of holonomies around these paths must be equal up to inversion; since the relative orientations of $\mathfrak{p}_{\mu_j}$ and $\mathfrak{p}'_{\mu_j}$ with respect to $\mathcal{C}_{\rm small}$ are reversed and the sheets of the cover Σ are *also* reversed at opposite ends of a defect, the eigenvalues $x_{\mathfrak{p}_{\mu_j}}$ and $x_{\mathfrak{p}'_{\mu_j}}$ are exactly equal.

Part (c) follows immediately from part (b) and Lemma 3.

Injectivity of (4.11) (and the identification of the image) in (d) follows from Appendix B, where we recreate the unique framed flat connection $\mathcal{A} \in \mathcal{P}[\mathcal{C}]$ associated to any point $x_\gamma \in (\mathbb{C}^*)^{d(\mathcal{X}[\mathcal{C}])}\big|_{x_\lambda^{(a)} \neq 1,\, (x_\alpha^{(t)}, x_\beta^{(t)}) \neq (1,1)}$.

The equality $\operatorname{rank} H_1^-(\Sigma) = d(\mathcal{X}[\mathcal{C}])$ in part (e) follows by comparing (2.10) and (4.3). □

When M has no defects and $\pi_1(\mathcal{C}_{\rm small})$ is abelian, the functions x_γ are actually quite familiar. Above, we already produced functions $x_\alpha^{(t)}, x_\beta^{(t)}$ for small tori and "lengths" $x_\lambda^{(a)}$ for small annuli. Going further, we may choose a path $\mathfrak{p}_\tau^{(a)}$ traversing each annulus. Then

$$x_\tau^{(a)} = x_{\mathfrak{p}_\tau^{(a)}} \tag{4.12}$$

is a complex generalization of the Fenchel-Nielsen "twist" function of Teichmüller theory. Finally, for every big edge e, we let $\mathfrak{p}_e, \mathfrak{p}'_e$ be the two paths running counter-clockwise around endpoints of e on the small boundary (Figure 20, right). It is easy to see that the functions $x_{\mathfrak{p}_e} = x_{\mathfrak{p}'_e}$ both reduce to a standard cross-ratio of the four framing lines surrounding the edge e. With orientation as on the right of Figure 20, we have

$$x_e := x_{\gamma_e} = x_{\mathfrak{p}(e)} = x_{\mathfrak{p}'(e)} = -\frac{\langle a \wedge b \rangle \langle c \wedge d \rangle}{\langle a \wedge c \rangle \langle b \wedge d \rangle}\,. \tag{4.13}$$

This is a standard Fock-Goncharov cross-ratio function[13], a generalization of Thurston's shear coordinates in Teichmüller theory.

The standard relations among functions x_e, x_τ, x_λ, arising from $x_{\mathfrak{p}\circ\mathfrak{p}'} = x_{\mathfrak{p}} x_{\mathfrak{p}'}$, are illustrated in Figure 21. Moving the endpoints of a path $\mathfrak{p}_\tau$ results in multiplication of x_τ by x_e functions; the product of x_e's around any hole filled in by a small disc equals 1 (consistent with the fact that holonomy around a small disc is trivial); and the product of x_e's around the two ends of an annulus are inverses of each other, and equal to $x_\lambda^{\pm 1}$. From injectivity of (4.11) and the presentation (3.10) of $H_1^-(\Sigma)$ in Section 3, it follows that these are the only

[13] This is the negative of the edge coordinate discussed in Appendix B of [32], and agrees with the positive coordinates of [44].

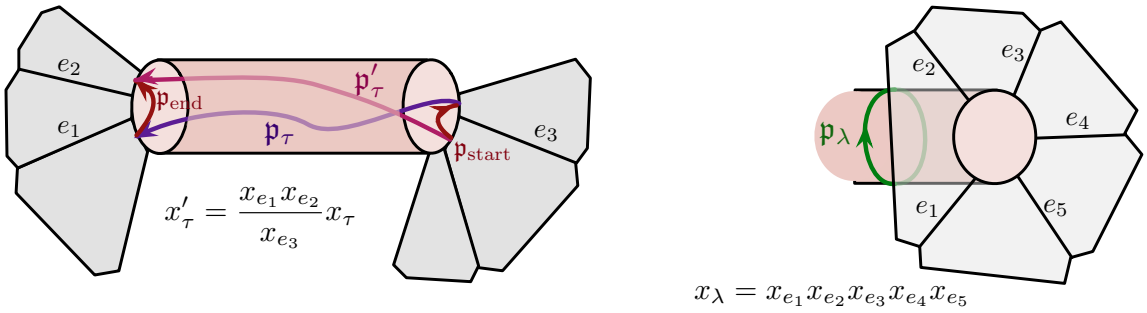

Figure 21. Examples of the basic relations among x_λ, x_τ, x_e: moving the endpoints of τ path (left); and expressing a holonomy as a product of edges (right).

non-trivial relations. Indeed, we must have

$$\mathcal{P}[\mathcal{C}] = \Big\{ (x_\alpha^{(t)}, x_\beta^{(t)}, x_\lambda^{(a)}, x_\tau^{(a)}, x_e) \in [(\mathbb{C}^*)^2 \backslash (1,1)]^{\#t} \times [(\mathbb{C}^* \backslash (1)) \times \mathbb{C}^*]^{\#a} \times (\mathbb{C}^*)^{\#e} \Big| \prod_{e \text{ on } d} x_e = 1, \prod_{e \text{ on } \partial_1 a} x_e = \prod_{e \text{ on } \partial_2 a} x_e^{-1} = \lambda^{(a)} \Big\}. \tag{4.14}$$

4.3 Poisson brackets, symplectic form, and K_2 form

When M has no defects and $\pi_1(\mathcal{C}_{\rm small})$ is abelian, we can calculate the Poisson bracket among functions x_γ at generic points of $\mathcal{X}[\mathcal{C}]$ by pulling-back the Atiyah-Bott formula (4.4). In fact, all the relevant calculations have already been done in [44] and [4, Appendix B]. (In both references, the idea was to calculate fundamental brackets involving the contractions of lines $\langle a \wedge b \rangle$ that appear in formulas such as (4.8) and (4.13). Holonomy eigenvalues x_λ, etc., can also easily be written in terms of such contractions.) We find

$$\begin{aligned}
\text{for each small torus:} \quad & \{\log x_\alpha, \log x_\beta\} = 2\langle \mathfrak{p}_\alpha, \mathfrak{p}_\beta \rangle = 2 \\
\text{for each small annulus:} \quad & \{\log x_\tau, \log x_\lambda\} = 2\langle \mathfrak{p}_\tau, \mathfrak{p}_\lambda \rangle = \pm 2 \\
& \{\log x_\tau, \log x_e\} = \pm 1 \qquad \text{if } e \text{ on faces adjacent to } \partial \mathfrak{p}_\tau \\
\text{for big edges } e, e' : \quad & \{\log x_e, \log x_{e'}\} = \#\text{ faces shared by } e, e' ,
\end{aligned} \tag{4.15}$$

with all other brackets vanishing. Here $\langle \mathfrak{p}_\tau, \mathfrak{p}_\lambda \rangle$ denotes an intersection number of paths on an annulus (similarly on a torus), with orientation such that $\langle \mathfrak{p}_\tau, \mathfrak{p}_\lambda \rangle = +1$ on the left of Figure 20 (the paths there intersect on the back side of the annulus). For the $\{\log x_\tau, \log x_e\}$ bracket, there are three edges bounding the face f at the start of $\mathfrak{p}_\tau$) and three bounding f' the end of $\mathfrak{p}_\tau$, with signs of brackets given in Figure 22. Finally, the number of faces in an edge-edge bracket is counted with orientation, such that $\{\log x_{e_1}, \log x_{e_2}\} = +1$ in Figure 22.

Although the brackets (4.15) are calculated at points of $\mathcal{P}[\mathcal{C}] \subset \mathcal{X}[\mathcal{C}]$ that project to smooth points of the standard character variety, we will use (4.15) to define them on the entire $(\mathbb{C}^*)^d$ patch $\mathcal{P}[\mathcal{C}]$. One then ought to verify that this extension is natural and consistent — in particular that it agrees on overlaps of patches $\mathcal{P}[\mathcal{C}; \mathbf{t}_{2\rm d}]$ labelled by *different* big-boundary triangulations. The main ingredient of such a verification (as in [44]) is the fact that a flip of

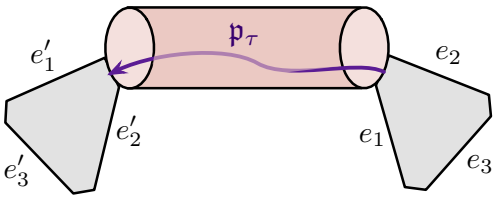

$$\{\log x_\tau, \log x_1\} = \{\log x_\tau, \log x_2\} = -\{\log x_\tau, \log x_3\} = 1$$
$$\{\log x_\tau, \log x_1'\} = \{\log x_\tau, \log x_2'\} = -\{\log x_\tau, \log x_3'\} = -1$$
$$\{\log x_1, \log x_2\} = \{\log x_2, \log x_3\} = \{\log x_3, \log x_1\} = 1$$

Figure 22. Orientation convention for Poisson brackets of twist and edge coordinates. In the formulas, we abbreviate $x_{e_i} = x_i$, $x_{e_i'} = x_i'$, etc.

2d triangulation acts as a symplectomorphism — in fact, a K_2 symplectomorphism. We will not pursue the global nature of Poisson/symplectic structures on framed moduli spaces here. Strictly speaking we do not need it: we will confine ourselves to working with fixed patches and triangulations.

By direct comparison to (3.11) in Lemma 4, we see that the Poisson brackets on $\mathcal{P}[\mathcal{C}]$ are all encoded in the intersection form in odd homology: $\{\log x_\gamma, \log x_{\gamma'}\} = \langle \gamma, \gamma' \rangle$ for any $\gamma, \gamma' \in H_1^-(\Sigma)$. The symplectic form on $\mathcal{P}[\mathcal{C}]$ can be expressed concretely by inverting the Poisson brackets,

$$\omega|_{\mathcal{P}[\mathcal{C}]} = \frac{1}{2} \sum_{i,j} (\epsilon^{-1})^{ij} \frac{dx_i}{x_i} \wedge \frac{dx_j}{x_j}\,, \qquad \epsilon_{ij} := \langle \gamma_i, \gamma_j \rangle\,, \tag{4.16}$$

where $\{\gamma_i\}$ is any basis for $H_1^-(\Sigma)$ and the $x_i := x_{\gamma_i}$ are corresponding coordinates on $\mathcal{P}[\mathcal{C}]$. (Remember that the intersection form on odd homology is non-degenerate, so ϵ is invertible, though not necessarily over the integers.) Moreover, the symplectic form lifts to a K_2 avatar

$$\hat{\omega} := \frac{1}{2} \sum_{ij} (\epsilon^{-1})^{ij} x_i \wedge x_j \quad \in K_2(\mathcal{P}[\mathcal{C}])_{\mathbb{Q}} \tag{4.17}$$

in the second algebraic K-group of the field of functions on $\mathcal{P}[\mathcal{C}]$ (tensored[14] with $\mathbb{Q}$). Such K_2 avatars are discussed from various perspectives in [13, 14, 31, 44]. We summarize these observations as:

Proposition 3 *If M has no defects and $\pi_1(\mathcal{C})$ is abelian, so that the functions x_γ ($\gamma \in H_1^-(\Sigma) \subset \widetilde{H}_1^-(\Sigma)$) contain complete coordinates on $\mathcal{P}[\mathcal{C}]$, the holomorphic symplectic form (agreeing with the Atiyah-Bott form), its K_2 avatar, and the holomorphic Poisson brackets are all encoded by the intersection product on $H_1^-(\Sigma)$. In particular,*

$$\{x_\gamma, x_{\gamma'}\} = \langle \gamma, \gamma' \rangle\, x_\gamma x_{\gamma'} \qquad \textit{or} \qquad \{\log x_\gamma, \log x_{\gamma'}\} = \langle \gamma, \gamma' \rangle\,. \tag{4.18}$$

[14] It follows from Lemma 7 (Appendix A.1) that ϵ^{-1} contains either integers or half-integers. Thus one could actually tensor with $[\frac{1}{2}]$ (killing 2^p torsion) rather than with $\mathbb{Q}$. We don't keep careful track of this here.

4.3.1 Example: tetrahedron

The punctured boundary of a tetrahedron $\mathcal{C}^*_\Delta$ is a four-punctured sphere, and the space of framed flat $PGL(2)$ connections on $\mathcal{C}^*_\Delta$ is two-dimensional. It has a canonical open subset $\mathcal{P}[\mathcal{C}_\Delta] = \mathcal{P}[\mathcal{C}_\Delta, \mathbf{t}_{2\mathrm{d}}] \simeq \mathbb{C}^* \times \mathbb{C}^* \subset \mathcal{X}[\Delta]$ corresponding to the canonical triangulation $\mathbf{t}_{2d}$ of the tetrahedron's big boundary [31]. Given the description of odd homology $H_1^-(\Sigma_\Delta)$ in Section 3.2.1, Proposition 2 says that $\mathcal{P}[\mathcal{C}_\Delta]$ is covered by six edge functions, subject to the relation that the product around every vertex is one. In terms of cross-ratios of framing lines, we have

$$x_e = -\frac{\langle a \wedge b\rangle\langle c \wedge d\rangle}{\langle a \wedge c\rangle\langle b \wedge d\rangle}\,, \quad x'_e = -\frac{\langle b \wedge d\rangle\langle c \wedge a\rangle}{\langle b \wedge c\rangle\langle d \wedge a\rangle}\,, \quad x''_e = -\frac{\langle d \wedge a\rangle\langle c \wedge b\rangle}{\langle d \wedge c\rangle\langle a \wedge b\rangle}\,, \tag{4.19}$$

with cross-ratios on opposite edges equal and $x_e x'_e x''_e = 1$ (Figure 23). Here $x_e = x_\gamma$, $x'_e = x_{\gamma'}$, etc. for cycles $\gamma, \gamma' \in \widetilde{H}_1^-(\Sigma_\Delta)$. The Poisson bracket and K_2 form in Prop. 3 are

$$\{\log x_e, \log x'_e\} = \{\log x'_e, \log x''_e\} = \{\log x''_e, \log x_e\} = 1\,, \qquad \hat{\omega} = x_e \wedge x'_e\,. \tag{4.20}$$

Setting

$$z = -x_e = x_{\gamma+u}\,, \quad z' = -x'_e, \quad z'' = -x''_e\,, \tag{4.21}$$

we recover the familiar phase space implicit in hyperbolic geometry [17].

In hyperbolic geometry there is usually an additional relation $z + z'^{-1} - 1 = 0$. This arises when considering framed flat connections on $\mathcal{C}_\Delta$ that can be extended to framed flat connections in the interior of the tetrahedron. Indeed, in this case the flat connection can be trivialized, and all three cross-ratios in (4.19) are computed in a single common fiber, leading to the additional relation. The submanifold $\mathcal{L}_\Delta = \{z + z'^{-1} - 1 = 0\} \subset \mathcal{P}[\mathcal{C}_\Delta]$ is a K_2 Lagrangian submanifold; it simply parameterizes the configuration space of four lines in $\mathbb{C}^2$. For further discussion, see [10, 17, 31].

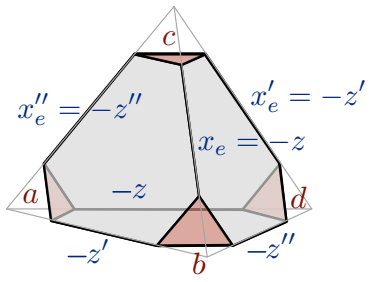

Figure 23. Functions on $\mathcal{X}[\mathcal{C}_\Delta]$

4.4 Gluing $PGL(2)$ connections on boundaries

Our next task is to describe the gluing equations for framed $PGL(2)$ flat connections on boundaries. To do so, we first explain abstractly what it means to glue framed flat connections on boundaries.

Suppose that we glue framed 3-manifolds $M \rightsquigarrow M_0 \rightsquigarrow M'$ in two steps, as in Section 2.5. Assume that M and M' are defect-free, and that their respective small boundaries $\mathcal{C}_{\mathrm{small}}$, $\mathcal{C}'_{\mathrm{small}}$ have abelian fundamental groups.

There is always a map

$$g^{(1)}_{PGL(2)} \,:\, \mathcal{X}[\mathcal{C}] \to \mathcal{X}[\mathcal{C}^{**}_0]\,, \tag{4.22}$$

Namely, given a framed flat connection $\mathcal{A}$ on $\mathcal{C}^*$, we may use a gauge transformation to trivialize it in the interiors of faces of $\mathbf{t}_{2\mathrm{d}}$, and then (trivially) identify connection on pairs

of faces that are glued together. We also identify framing flags on adjacent pieces of small boundary. We must recall, however, that by definition $\mathcal{A}$ is a framed flat connection on the punctured boundary $\mathcal{C}^* = \mathcal{C}\backslash\{\text{point on each small disc}\}$ rather than on $\mathcal{C}$ itself, so gluing produces a flat connection not on $\mathcal{C}_0$ itself but on $\mathcal{C}_0$ with some collection of points (punctures) removed from its small boundary, which we denote $\mathcal{C}_0^{**}$. The connection is framed on $(\mathcal{C}_0^{**})_{\text{small}}$, and has unipotent holonomy around each puncture.

Next, in order to fill in the defects I_j of $\mathcal{C}_0^{**}$, it is necessary and sufficient to require that the holonomy M_j around each defect is trivial. (If the holonomy is trivial, then connection can be trivialized along the defect, which allows the defect to be filled in.) Therefore, we get a map

$$g^{(2)}_{PGL(2)} \,:\, \mathcal{X}[\mathcal{C}_0^{**}]\big|_{M_j=I} \to \mathcal{X}[\mathcal{C}'^{**}]\,. \tag{4.23}$$

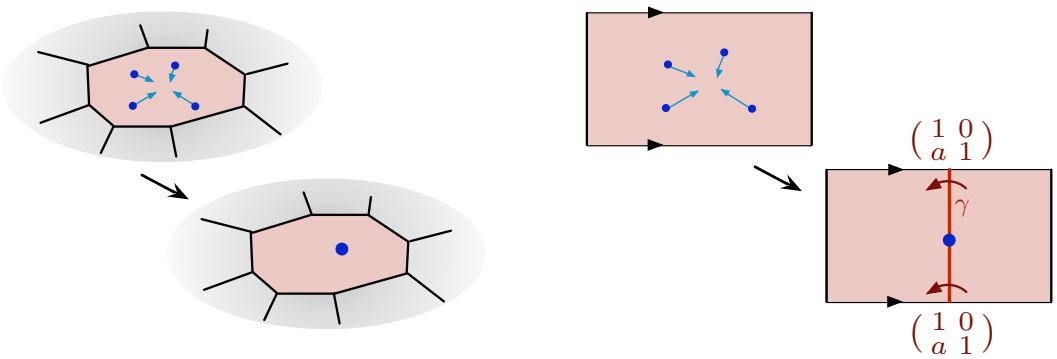

Figure 24. Removing extraneous punctures on small boundaries. For a disc (left) we just collapse the punctures. For an annulus or torus (right) we collapse the punctures and modify the flat connection to trivialize the unipotent holonomy around the puncture.

Again, we observe that $\mathcal{C}'^{**}$ may have extraneous punctures, arising from disjoint discs in $\mathcal{C}_{\text{small}}$ that have been connected in $\mathcal{C}'_{\text{small}}$. We can remove these punctures. For every small sphere component of $\mathcal{C}'_{\text{small}}$, we collapse all (potential) punctures to a single one; the holonomy around this single puncture is necessarily trivial, so we may remove it. For every small disc, we similarly collapse all (potential) punctures to a single one, which still has unipotent holonomy (Figure 24, left). For every small annulus we modify the connection so that it extends over the punctures. There is a unique way to do this. Abstractly, we note that since all puncture holonomies are unipotent and the connection is framed on the annulus (meaning, as usual, that it preserves a line) it makes sense to define an invariant holonomy eigenvalue x_λ and a rescaling/twist coordinate x_τ from one end of the annulus to the other, just as in Section 4.2. In fact, we can calculate full (basepointed) holonomy matrices M_λ, M'_λ, both with eigenvalues x_λ, around the two ends of the annulus. Then from Appendix B we know that there's a unique framed flat connection on the smooth annulus with fixed $M_\lambda, M'_\lambda, x_\tau$, as long as $x_\lambda \neq 1$.

To implement the modification concretely, we can collapse all punctures on an annulus to a single point, then cut the annulus along a curve γ beginning and ending at this point

as in Figure 24 (right). We re-glue the connection on the annulus with an extra unipotent modification $\left(\begin{smallmatrix} 1 & 0 \\ a & 1 \end{smallmatrix}\right)$ across γ. As long as $x_\lambda \neq 1$, there is a unique value of a that trivializes the puncture holonomy.

We may similarly remove extraneous punctures from small tori, so long as the invariant holonomy eigenvalues satisfy $(x_\alpha, x_\beta) \neq (1,1)$. Concretely, we collapse all punctures to a single one, then add a unipotent modification along a nontrivial curve γ passing through the single puncture. The modification that trivializes the puncture holonomy is uniquely specified so long as the holonomy along γ itself is not unipotent — so, given $(x_\alpha, x_\beta) \neq (1,1)$, an appropriate γ can always be chosen.

Altogether, we have built a removal-of-punctures map

$$g^{(3)}_{PGL(2)} \ : \ \mathcal{X}[\mathcal{C}'^{**}]\big|_{x_\lambda \neq 1,\, (x_\alpha, x_\beta) \neq (1,1)} \to \mathcal{X}[\mathcal{C}'] \,, \tag{4.24}$$

and composing it with the first two parts of the gluing procedure we obtain a gluing map

$$\boxed{g_{PGL(2)} = g^{(3)}_{PGL(2)} \circ g^{(2)}_{PGL(2)} \circ g^{(1)}_{PGL(2)} \ : \ \mathcal{X}[\mathcal{C}]\big|_R \to \mathcal{X}[\mathcal{C}']} \,. \tag{4.25}$$

Here the restrictions 'R' are a lift to $\mathcal{X}[\mathcal{C}]$ of the defect conditions $M_j = 1$ and puncture-removal conditions $x_\lambda \neq 1$, $(x_\alpha, x_\beta) \neq (1,1)$.

4.5 Symplectic properties of the gluing equations

Let's now restrict to patches $\mathcal{P}[\mathcal{C}, \mathbf{t}_{2\mathrm{d}}]$, $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$, corresponding to the triangulation of M and the induced triangulation of M', and describe concretely how the gluing map (4.25) acts on $\mathbb{C}^*$ coordinates.

Consider the intermediate space $\mathcal{X}[\mathcal{C}_0^{**}]$, where $\mathcal{C}_0^{**}$ is the boundary of M_0, possibly with some additional unipotent punctures. This space supports the same path-functions $x_{\mathfrak{p}}$, $\mathfrak{p} \in \mathbf{P}_0 \oplus \mathbb{Z}_2$ as $\mathcal{X}[\mathcal{C}_0]$ (since unipotent punctures do not affect the definition or properties of the $x_{\mathfrak{p}}$), which depend only on the homology class $\gamma \in \widetilde{K} = \mathrm{im}\left[\widetilde{h} : (\mathbf{P}_0 \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma)\right]$. As usual, we work with the convention that $x_u = -1$ for the fiber class $u \in \mathbb{Z}_2$. The functions x_γ, $\gamma \in \widetilde{K}$ take well-defined $\mathbb{C}^*$ values on a subset $\mathcal{P}[\mathcal{C}_0^{**}, \mathbf{t}'_{2\mathrm{d}}] \subset \mathcal{X}[\mathcal{C}_0^{**}]$ that is defined the same way as in Def. 3.

The image[15] of $g^{(1)}_{PGL} : \mathcal{P}[\mathcal{C}; \mathbf{t}_{2\mathrm{d}}] \to \mathcal{X}[\mathcal{C}_0^{**}]$ lies in $\mathcal{P}[\mathcal{C}_0^{**}; \mathbf{t}'_{2\mathrm{d}}]$. Moreover,

$$g^{(1)}_{PGL}{}^*(x_\gamma) = x_{\widetilde{g}(\gamma)} \,, \qquad \gamma \in \widetilde{K} \,, \tag{4.26}$$

with $\widetilde{g} : \widetilde{H}_1^-(\Sigma_0) \to \widetilde{H}_1^-(\Sigma)$ the map of (2.18) and (3.21). To understand this, suppose $\gamma = \widetilde{h}(\mathfrak{p})$ for a path $\mathfrak{p} \in \mathbf{P}_0$. We can compute the path coordinate $x_\gamma = x_{\mathfrak{p}}$ by breaking $\mathfrak{p}$ up into segments $g_{\mathbf{P}}(\mathfrak{p}) = \sum_i \mathfrak{p}_i$, $\mathfrak{p}_i \in \mathbf{P}$ (as in Figure 18), applying the definition (4.8) to each segment to get functions $x_{\mathfrak{p}_i}$, and multiplying them together. Successive numerators and denominators of (4.8) cancel out — just as in the proof of the concatenation relation $x_{\mathfrak{p}\circ\mathfrak{p}'} =$

[15] Technically, we should relax the non-unipotent restriction in the definition of $\mathcal{P}[\mathcal{C}_0^{**}]$ in order for the image of $g^{(1)}_{PGL} : \mathcal{P}[\mathcal{C}] \to \mathcal{X}[\mathcal{C}_0^{**}]$ to be fully contained in $\mathcal{P}[\mathcal{C}_0^{**}]$. This does not affect the following argument.

$x_{\mathfrak{p}}x_{\mathfrak{p}'}$ — up to a single sign (-1). This extra sign arises due to the reversed relative orientation at the head of one segment and the tail of the next. Thus $x_{\mathfrak{p}} = ...x_{\mathfrak{p}_{i+1}}(-1)x_{\mathfrak{p}_i}(-1)x_{\mathfrak{p}_{i-1}}....$ The modifications by these signs are precisely encoded in the fiber-class corrections to the twisted cutting map $\widetilde{g}_{\mathbf{P}}$ (3.21). Using $\widetilde{h}\widetilde{g}_{\mathbf{P}} = \widetilde{g}\widetilde{h}$, we obtain (4.26).

The gluing conditions requiring trivial defect holonomies M_j simply say that $x_{\mu_j} \equiv 1$ for all defect cycles $\mu_j \in \widetilde{G} \subset \widetilde{K}$. The lifts of these conditions to $\mathcal{P}[\mathcal{C}, \mathbf{t}_{2\mathrm{d}}]$ are $x_{\tilde{g}(\mu_j)} \equiv 1$. If these conditions are satisfied, we get a map

$$g^{(3)}_{PGL(2)} \circ g^{(2)}_{PGL(2)} \,:\, \mathcal{P}[\mathcal{C}_0^{**}, \mathbf{t}'_{2\mathrm{d}}]\big|_{x_\mu \equiv 1} \to \mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]\,. \tag{4.27}$$

Now the functions x_γ ($\gamma \in \widetilde{H}_1^-(\Sigma')$) on $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$ pull back[16] to

$$\big(g^{(3)}_{PGL(2)} \circ g^{(2)}_{PGL(2)}\big)^*(x_\gamma) = x_{\hat{\gamma}}\,, \qquad \hat{\gamma} \in \widetilde{K}\,,\ \widetilde{q}(\hat{\gamma}) = \gamma\,, \tag{4.28}$$

where $\hat{\gamma} \in \widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)$ is any preimage of γ under the map $\widetilde{q}$ from (3.17). This preimage is unique modulo $\widetilde{G}$, so $x_{\hat{\gamma}}$ is well defined on $\mathcal{P}[\mathcal{C}_0^{**}, \mathbf{t}'_{2\mathrm{d}}]\big|_{x_\mu \equiv 1}$.

By combining (4.26) and (4.28) we see that for any $\gamma \in \widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)$, the path-functions on $\mathcal{P}[\mathcal{C}, \mathbf{t}_{2\mathrm{d}}]\big|_{M_j=1}$ and $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$, evaluated on a connection $\mathcal{A}$ and its glued image $g_{PGL(2)}(\mathcal{A})$, are related by

$$\boxed{x_{\tilde{g}(\gamma)} = x_{\tilde{q}(\gamma)}\,, \qquad \forall\, \gamma \in \widetilde{K} \subset \widetilde{H}_1^-(\Sigma_0)\,.} \tag{4.29}$$

These are the gluing equations. They subsume the gluing conditions $x_{\tilde{g}(\mu)} \equiv 1$ for $\mu \in \widetilde{G}$, since $\widetilde{q}(\mu) \equiv 0$. Moreover, since $\widetilde{q}$ is surjective and the path-functions on $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$ include a complete set of coordinates, the gluing map

$$g_{PGL(2)} \,:\, \mathcal{P}[\mathcal{C}, \mathbf{t}_{2\mathrm{d}}]\big|(x_{\tilde{g}(\mu)} \equiv 1, R') \twoheadrightarrow \mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}] \tag{4.30}$$

must be surjective. (Here R' denotes the additional non-unipotent restrictions lifted from $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$, namely $x_{\tilde{g}(\hat{\lambda})} \neq 1$, $(x_{\tilde{g}(\hat{\alpha})}, x_{\tilde{g}(\hat{\beta})}) \neq (1,1)$, for $\widetilde{q}(\hat{\lambda}, \hat{\alpha}, \hat{\beta}) = (\lambda, \alpha, \beta)$.) Indeed,

Theorem 1 *The $PGL(2)$ gluing map (4.30) is a (holomorphic) K_2 symplectic reduction of a finite quotient,*

$$g_{PGL(2)} \,:\, (\mathcal{P}[\mathcal{C}; \mathbf{t}_{2\mathrm{d}}]_{R'}/Z)/\!\!/(\mathbb{C}^*)^{\mathrm{rank}\,G} := (\mathcal{P}[\mathcal{C}; \mathbf{t}_{2\mathrm{d}}]_{R'}/Z)\big|_{x_{\tilde{g}(\mu)} \equiv 1}/(\mathbb{C}^*)^{\mathrm{rank}\,G} \\ \tilde{\to}\ \mathcal{P}[\mathcal{C}'; \mathbf{t}'_{2\mathrm{d}}]\,, \tag{4.31}$$

where $Z \simeq \widetilde{H}_1^-(\Sigma)/\widetilde{H}$ is the torsion group of Prop. 1'(b), and the $(\mathbb{C}^*)^{\mathrm{rank}\,G}$ action is generated by using the $x_{\tilde{g}(\mu_j)}$, $\mu_j \in \widetilde{G}$, as moment maps.

[16] Again, note that the unipotent modifications made by $g^{(3)}_{PGL(2)}$ to remove extraneous punctures do not affect path coordinates. So to understand (4.28) it suffices to look at the action of $g^{(2)}_{PGL(2)}$, which just fills in defects.

Explicitly, if $\widetilde{H} \subset \widetilde{H}_1^-(\Sigma)$ is the finite-index subgroup containing $g(\widetilde{K})$, such that $\widetilde{H}_1^-(\Sigma') \simeq \widetilde{H}//\widetilde{g}(\widetilde{G}) = \widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G})$, we recall that $Z = \widetilde{H}_1^-(\Sigma)/\widetilde{H}$ contains at most 2-torsion and 4-torsion. The group Z act naturally on $\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]$ by multiplying functions x_γ by 4-th roots of unity. Namely, letting $\{\sigma_i\}$ be any generators of the Hom $\big(Z, \mathbb{Z}_4 = \{1, i, -1, -i\}\big)$, the action on $\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]$ is generated by $x_\gamma \mapsto \sigma_i[\gamma] x_\gamma$.

We may also explicitly describe the action of the moment maps as follows. Let $\{\mu_i\}_{i=1}^{\operatorname{rank} G}$ be a basis for $\widetilde{G}$ and let $(t_i)_{i=1}^{\operatorname{rank} G} \in (\mathbb{C}^*)^{\operatorname{rank} G}$ be some corresponding parameters. If $\mu = \sum a_i \mu_i \in \widetilde{G}$ define $t^\mu = \prod_i t_i^{a_i}$. Then, by virtue of (4.18) in Proposition 3, the moment map action of $x_{\widetilde{g}(\mu)}$ is

$$x_\gamma \mapsto (t^\mu)^{\langle \mu, \gamma \rangle} x_\gamma \,, \qquad \forall\, \gamma \in \widetilde{H}_1^-(\Sigma)\,,\ \mu \in \widetilde{G}\,. \tag{4.32}$$

Proof of Theorem 1. We analyze the quotient on the LHS of (4.31). By construction, the invariant functions on the ordinary quotient $\mathcal{P}[\mathcal{C},\mathbf{t}_{2\mathrm{d}}]/Z$ are precisely x_γ for $\gamma \in \widetilde{H} \subset \widetilde{H}_1^-(\Sigma)$. These include a complete set of coordinates. Since $g(\widetilde{K}) = \ker \langle \widetilde{g}(\widetilde{G}), * \rangle|_{\widetilde{H}}$ by Proposition 1' and Lemma 5, the functions that are further invariant under the $(\mathbb{C}^*)^{\operatorname{rank} G}$ action (4.32) are labeled by $\gamma \in \widetilde{g}(\widetilde{K}) \subset \widetilde{H}_1^-(\Sigma)$. These functions coordinatize $(\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]/Z)/(\mathbb{C}^*)^{\operatorname{rank} G}$. Subsequently restricting to $x_{\widetilde{g}(\mu)} \equiv 1$ for all $\mu \in \widetilde{G}$ produces a space coordinatized by $x_{[\gamma]}$ for $[\gamma] \in \widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G}) \simeq \widetilde{H}_1^-(\Sigma')$. The identification $\widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G}) \simeq \widetilde{H}_1^-(\Sigma')$ sends $[\gamma]$ to $\widetilde{q} \circ \widetilde{g}^{-1}(\gamma)$ and preserves the intersection form. Therefore, $(\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]_{R'}/Z)/(\mathbb{C}^*)^{\operatorname{rank} G}\big|_{x_{\widetilde{g}(\mu)}\equiv 1}$ is canonically (1-1) symplectomorphic to $\mathcal{P}[\mathcal{C}';\mathbf{t}'_{2\mathrm{d}}]$, with path-coordinates related as $x_{\widetilde{g}(\gamma)} = x_{\widetilde{q}(\gamma)}$ for $\gamma \in \widetilde{K}$. These are precisely the gluing equations (4.30), so the symplectomorphism is the desired gluing map.

To see that this is a K_2 symplectomorphism, first note that $\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]$ and $\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]/Z$ have the same K_2 forms $\hat{\omega} \in K_2(\mathcal{P}[\mathcal{C}])_{\mathbb{Q}}$, given by (4.17). Let us choose a basis $\{\gamma_i, \mu_j\}$ for the untwisted subgroup $K \subset \widetilde{K}$ such that $\{\widetilde{q}(\gamma_i)\}$ is a basis for $H_1^-(\Sigma') \subset \widetilde{H}_1^-(\Sigma')$. Complete this to a basis $\{\widetilde{g}(\gamma_i), \widetilde{g}(\mu_j), \eta_k\}$ of a torsion-free (untwisted) subgroup $H \subset \widetilde{H}$, so that the intersection form in this basis is block-diagonal,

$$\langle \widetilde{g}(\gamma_i), \widetilde{g}(\gamma_{i'}) \rangle = \langle \widetilde{q}(\gamma_i), \widetilde{q}(\gamma_{i'}) \rangle = \epsilon'_{ii'}\,, \qquad \langle \widetilde{g}(\mu_j), \eta_k \rangle = c_j\, \delta_{jk} \quad (c_j \in \mathbb{Z})\,, \tag{4.33}$$

where ϵ' is the intersection form on $H_1(\Sigma')$ and all other intersection products vanish. (The fact that we can find such a basis follows from the untwisted version of $\widetilde{H}//\widetilde{g}(\widetilde{G}) = \widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G}) \simeq \widetilde{H}_1^-(\Sigma')$.) Then the K_2 form on $\mathcal{P}[\mathcal{C};\mathbf{t}_{2\mathrm{d}}]/Z$ is

$$\hat{\omega} = \frac{1}{2} \sum_{i,i'} (\epsilon'^{-1})^{ii'} x_{\widetilde{g}(\gamma_i)} \wedge x_{\widetilde{g}(\gamma'_i)} + \sum_j \frac{1}{c_j} x_{\widetilde{g}(\mu_j)} \wedge x_{\eta_j}\,. \tag{4.34}$$

Setting $x_{\widetilde{g}(\mu_j)} = 1$ and using the fact that $1 \wedge * = 0$ in K-theory, we find that $\hat{\omega}|_{(x_{\widetilde{g}(\mu)} \equiv 1)} = \hat{\omega}'$ reproduces the K_2 form on $\mathcal{P}[\mathcal{C}', \mathbf{t}'_{2\mathrm{d}}]$. □

4.6 Example: Thurston's gluing equations, K_2 forms, and volumes

In the Introduction (Section 1.3), we claimed that Thurston's gluing equations are a special case of (4.29), and therefore that Theorem 1 directly implies the symplectic properties found by Neumann and Zagier. We take a moment to spell out exactly how this works.

Suppose that M' is an oriented hyperbolic 3-manifold with an ideal triangulation $M' = \cup_{i=1}^N \Delta_i$. We denote by $M = \sqcup_{i=1}^N \Delta_i$ the disjoint union of tetrahedra in the triangulation, viewing both M and M' as framed 3-manifolds. For example, M' might be a closed hyperbolic manifold with a spun triangulation, in which case the topological ideal boundary $\mathcal{C}' = \partial M'$ consists entirely of small spheres; or M' could have n_c cusps, corresponding to small torus components of $\mathcal{C}'$. One might also consider M' with geodesic boundary, in which case $\mathcal{C}'_{\text{big}}$ is non-empty. This is another special case of our general framework, but we'll ignore it for the moment to keep this example simple. Thus, $\mathcal{C}'_{\text{big}}$ is empty, and $\mathcal{C}'_{\text{small}}$ contains $n_c \geq 0$ small tori and some number of small spheres. The number of edges I_j of the triangulation is the same as the number N of tetrahedra. We glue $M \rightsquigarrow M_0 \rightsquigarrow M'$ as in Section 2.5, letting M_0 be a framed 3-manifold with N defects I_j.

For each tetrahedron Δ_i in M we find a K_2 symplectic space $\mathcal{P}[\mathcal{C}_{\Delta i}] = \{x_i, x'_i, x''_i \,|\, x_i x'_i x''_i = 1\} \simeq (\mathbb{C}^*)^2$ as in Section 4.3.1. Here $x_i = x_{\gamma_i}, x'_i = x_{\gamma'_i}$, etc., with $\gamma_i, \gamma'_i, \gamma''_i$ and u generating $\widetilde{H}_1^-(\Sigma_\Delta)$ as in Section 3.2.1. The triple of hyperbolic shapes for each tetrahedron are $(z_i, z'_i, z''_i) = (-x_i, -x'_i, -x''_i)$. Altogether,

$$\mathcal{P}[\mathcal{C}] = \prod_{i=1}^N \mathcal{P}[\mathcal{C}_{\Delta i}] \simeq (\mathbb{C}^*)^{2N}\,, \qquad \widetilde{H}_1^-(\Sigma) = \oplus_{i=1}^N H_1^-(\Sigma_\Delta) \oplus \mathbb{Z}_2 \simeq \mathbb{Z}^{2N} \oplus \mathbb{Z}_2\,, \tag{4.35}$$

where the splitting of $\widetilde{H}_1^-(\Sigma)$ comes naturally from Lemma 3. Similarly, for each small torus T_t^2 in $\mathcal{C}'_{\text{small}}$, there is an odd homology group $\widetilde{H}_1^-(\Sigma_{T^2})$ generated by $(\alpha^{(t)}, \beta^{(t)}, u)$ and a moduli space $\mathcal{P}[\mathcal{C}_{T^2}] = \{x_\alpha^{(t)}, x_\beta^{(t)}\} \simeq (\mathbb{C}^*)^2 \backslash (1,1)$. Thus

$$\mathcal{P}[\mathcal{C}'] = \prod_{t=1}^{n_c} \mathcal{P}[\mathcal{C}_{T_t^2}] \simeq (\mathbb{C}^*)^{2n_c}\,, \qquad \widetilde{H}_1^-(\Sigma') = \oplus_{i=1}^N H_1^-(\Sigma_{T_t^2}) \oplus \mathbb{Z}_2 \simeq \mathbb{Z}^{2n_c} \oplus \mathbb{Z}_2\,. \tag{4.36}$$

Now consider the intermediate manifold M_0. The boundary $\mathcal{C}_0 = \partial M_0$ small and defect parts; N defects end at $2N$ holes on $(\mathcal{C}_0)_{\text{small}}$. The holes on $(\mathcal{C}_0)_{\text{small}}$ lie at the vertices of the triangular tiling of $\mathcal{C}'_{\text{small}}$, as illustrated back in 3. The subgroup $\widetilde{G} = \text{im}\,[\widetilde{h} : \mathbf{P}_G \to \widetilde{H}_1^-(\Sigma_0)]$ is generated by cycles μ_j (Figure 16), one for each defect I_j, coming from paths $\mathfrak{p}_{\mu_j}$ that surround the holes of $(\mathcal{C}_0)_{\text{small}}$. The subgroup $\widetilde{K} = \text{im}\,[\widetilde{h} : (\mathbf{P}_0 \oplus \mathbb{Z}_2) \to \widetilde{H}_1^-(\Sigma_0)]$ is generated by the μ_j, together with lifts $\hat{\alpha}^{(t)}, \hat{\beta}^{(t)}$ (Section 3.3) of $\alpha^{(t)}, \beta^{(t)} \in H_1^-(\Sigma')$, and the fiber class u. Concretely, $\hat{\alpha}^{(t)} = \widetilde{h}(\hat{\mathfrak{p}}_\alpha^{(t)})$ and $\hat{\beta}^{(t)} = \widetilde{h}(\hat{\mathfrak{p}}_\beta^{(t)})$, where $\hat{\mathfrak{p}}_\alpha^{(t)}, \hat{\mathfrak{p}}_\beta^{(t)}$ are paths on $\mathcal{C}_{\text{small}}$ representing A and B cycles, passing in some chosen way around the holes. These are exactly the types of boundary paths that appeared in [12, 29], said to be in normal position with respect to the tiling of $\mathcal{C}'_{\text{small}}$.

The first part of Thurston's gluing equations state that the product of tetrahedron shapes z_i, z'_i, z''_i around any edge I_j must equal one. These are precisely our trivial-holonomy constraints, of the form

$$x_{\tilde{g}(\mu_j)} = 1\,, \qquad \forall\, \mu_j \in \widetilde{G}\,. \tag{4.37}$$

To see this, note that under the cutting map $g_{\mathbf{P}}$ (Section 3.3) a path $\mathfrak{p}_{\mu_j}$ surrounding defect I_j is cut into paths $\mathfrak{p}_e$ associated to all the edges e of tetrahedra identified with I_j in the

gluing. The extra modification in the twisted map $\widetilde{g}_{\mathbf{P}}$ adds a fiber class u to $\widetilde{g}_{\mathbf{P}}(\mathfrak{p}_{\mu_j})$ for every cut that is made; thus from $\widetilde{h}\widetilde{g}_{\mathbf{P}} = \widetilde{g}\widetilde{h}$ we find

$$\widetilde{g}(\mu_j) = \sum_{\text{around } I_j} (\gamma_i + u \text{ or } \gamma_i' + u \text{ or } \gamma_i'' + u)\,,$$

$$x_{\widetilde{g}(\mu_j)} = \prod_{\text{around } I_j} (-x_i, -x_i', -x_i'') = \prod_{\text{around } I_j} (z_i, z_i', z_i'')\,, \tag{4.38}$$

matching the hyperbolic gluing equations.

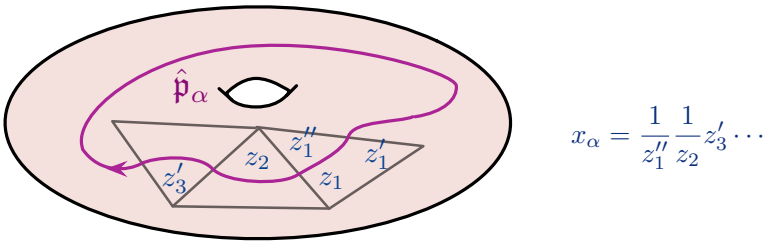

Figure 25. Calculating the cusp equation for $\ell^2 = x_\alpha^{(t)}$ on a small boundary that's tiled by small vertex triangles of truncated tetrahedra.

The second part of the gluing equations, sometimes called the "cusp equations," states that the squares of A and B cycle eigenvalues ℓ_t^2, m_t^2 (*i.e.* eigenvalues of the hyperbolic holonomy) around each cusp equal the product of shapes $z_i^{\pm 1}, z_i'^{\pm 1}, z_i''^{\pm 1}$ at dihedral angles subtended by paths on $\mathcal{C}_{\text{small}}$, representing the respective cycles and in normal position with respect to the tiling of $\mathcal{C}_{\text{small}}$ — in other words, our paths $\hat{\mathfrak{p}}_\alpha^{(t)}$ and $\hat{\mathfrak{p}}_\beta^{(t)}$. The exponents ± 1 correspond to whether an angle is subtended clockwise or counterclockwise, *cf.* Figure 25. These cusp equations again are just the remaining equations of the form $x_{\tilde{q}(\gamma)} = x_{\tilde{g}(\gamma)}$ in (4.29), for $\gamma = \hat{\alpha}^{(t)}$ or $\gamma = \hat{\alpha}^{(t)}$. Indeed, on one hand,

$$\ell_t^2 = x_\alpha^{(t)} = x_{\tilde{q}(\hat{\alpha}^{(t)})}\,, \qquad m_t^2 = x_\beta^{(t)} = x_{\tilde{q}(\hat{\beta}^{(t)})}\,. \tag{4.39}$$

On the other hand $\tilde{g}(\hat{\alpha}^{(t)}) = \widetilde{h} \circ \widetilde{g}_{\mathbf{P}}(\hat{\mathfrak{p}}_\alpha^{(t)})$ is a sum of cycles $\pm(\gamma_i + u)$, $\pm(\gamma_i' + u)$, $\pm(\gamma_i'' + u)$ corresponding to angles subtended by $\hat{\mathfrak{p}}_\alpha^{(t)}$ (similarly for $\hat{\beta}^{(t)}$); so $x_{\tilde{g}(\hat{\alpha}^{(t)})}$ and $x_{\tilde{g}(\hat{\beta}^{(t)})}$ are precisely the desired products of shapes.

Now Theorem 1 amounts to the statement that the defect functions $x_{\tilde{g}(\mu_j)}$ are a set of mutually commuting moment maps on $\mathcal{P}[\mathcal{C}]$, which also commute with the A and B cycle functions $x_{\tilde{g}(\hat{\alpha}^{(t)})}$ and $x_{\tilde{g}(\hat{\beta}^{(t)})}$. Thus $\mathcal{P}[\mathcal{C}']$ is the symplectic reduction of (a finite quotient of) $\mathcal{P}[\mathcal{C}]$. Obviously this implies that there are $\frac{1}{2}(\dim_{\mathbb{C}} \mathcal{P}[\mathcal{C}] - \dim_{\mathbb{C}} \mathcal{P}[\mathcal{C}']) = N - n_c$ independent moment maps. Homologically, this follows from the easy fact that

$$\operatorname{rank} \widetilde{G} = \operatorname{rank} G = N - n_c\,. \tag{4.40}$$

Indeed, the sum of generators μ_j for all defects that end on a given cusp (counted with multiplicity) is null-homologous, so each cusp produces one relation among the N generators.[17]

Alternatively, one could write the gluing equations in a fixed basis as

$$\begin{aligned} \ell_t^2 &= \pm \textstyle\prod_i z_i^{A_{ti}} z_i'^{A'_{ti}} \\ m_t^2 &= \pm \textstyle\prod_i z_i^{B_{ti}} z_i'^{B'_{ti}} \\ 1 = x_{\tilde{q}(\mu_j)} &= \pm \textstyle\prod_i z_i^{C_{ji}} z_i'^{C'_{ji}} \end{aligned} , \qquad g \begin{pmatrix} \alpha \\ \beta \\ \mu \end{pmatrix} = \begin{pmatrix} A & A' \\ B & B' \\ C & C' \end{pmatrix} \begin{pmatrix} \gamma \\ \gamma' \end{pmatrix} . \tag{4.41}$$

Then the matrix of the untwisted map $g|_K = p \circ \widetilde{g}_K$ shown here has rank $\operatorname{rank} K = N + n_c$ and preserves the intersection form, *i.e.* $(g|_K) J_{2N} (g|_K)^T = 2J_{2n_c} \oplus 0_{N\times N}$, as in the Introduction.

The reduction of K_2 forms is expressed as

$$\hat{\omega}' \;=\; \sum_{t=1}^{n_c} \frac{1}{2}\, \ell_t^2 \wedge m_t^2 \;=\; \sum_{i=1}^{N} z_i \wedge z_i' \Big|_{(x_{\tilde{g}(\mu)} \equiv 0)} \;=\; \hat{\omega}\big|_{(x_{\tilde{g}(\mu)} \equiv 0)} . \tag{4.42}$$

There is a standard map $\eta : K_2(\mathcal{P}[\mathcal{C}']) \to \Omega^1(\mathcal{P}[\mathcal{C}'])$ defined by $\eta(a \wedge b) = \log|a|\, d\arg b - \log|b|\, d\arg a$ (see [31] and references therin), which provides a canonical anti-derivative of the symplectic form on $\mathcal{P}[\mathcal{C}']$, since $d\eta(\hat{\omega}') = \operatorname{Im} \omega'$. By computing $\eta(\hat{\omega}')$ using both sides of (4.42) and further restricting to the K_2 Lagrangian $\mathcal{L} \subset \mathcal{P}[\mathcal{C}']$ defined by $z_i + z_i^{-1} - 1 = 0$ together with the gluing equations,[18] we directly recover the Neumann-Zagier formula for variation of the volume,

$$d\operatorname{Vol}(M') = \eta(\hat{\omega}')\big|_{\mathcal{L}} . \tag{4.43}$$

5 Non-abelianization

We showed in Section 4 that coordinates for framed flat $PGL(2)$ connections on the boundaries of framed 3-manifolds can be labelled by elements of the homology of double covers, with Poisson brackets of the former matching the intersection product of the latter. Moreover, we found that gluing equations (4.29) have a homological interpretation that makes their symplectic properties manifest. However, we did not explain *why* the relation between $PGL(2)$ connections and homology of double covers existed, or why it was particularly natural. We now aim to fill this gap.

We will first consider the moduli space of flat $GL(1) := GL(1, \mathbb{C})$ connections on a double cover Σ. This very simple space naturally has holonomy coordinates labelled by the first homology of Σ (viewed as an abelianization of $\pi_1(\Sigma)$), with Atiyah-Bott Poisson bracket induced by the intersection form. The gluing equations for flat $GL(1)$ connections manifestly take the form (4.29).

Then we borrow and extend the non-abelianization construction of [6] to build a non-abelianization map Φ, a (nontrivial) symplectomorphism between flat $GL(1)$ flat connections

[17] It is a short exercise to show that these relations are all independent. See [29] or the review [48].

[18] This "gluing variety" generalizes the geometric component of the A-polynomial for a knot [12–14].

on a double cover Σ and framed flat $PGL(2)$ connections on the base $\mathcal{C}$. While the non-abelianization map of [6] was mainly discussed in a rich geometric context — involving a choice of complex structure on $\mathcal{C}$ and an interpretation of Σ as a spectral cover — we will simply used the topological structure of the boundary of a framed 3-manifold (and a choice of big-boundary triangulation) to define Φ.[19]

5.1 Abelian flat connections

We begin by defining a space of abelian flat connections whose coordinates are manifestly labelled by elements of twisted homology $\widetilde{H}_1^-(\Sigma)$. As observed in [6] (and hinted in [44]), the non-abelianization construction requires such twisting.

For a closed oriented surface Σ, define

$$\begin{aligned}\widetilde{\mathcal{X}}_{\rm ab}[\Sigma] &= \{\text{twisted } GL(1) \text{ flat connections on } \Sigma\} \\ &:= \{GL(1) \text{ flat connections on } T_1\Sigma \text{ with fiber holonomy } -1\}\,.\end{aligned} \tag{5.1}$$

These are flat connections on a (necessarily trivial) complex line bundle $L \to T_1\Sigma$, or, equivalently, $GL(1) := GL(1,\mathbb{C})$ local systems on $T_1\Sigma$. The space $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma]$ is isomorphic to $(\mathbb{C}^*)^{\operatorname{rank}\tilde{H}_1(\Sigma)} = (\mathbb{C}^*)^{\operatorname{rank}H_1(\Sigma)}$. Indeed, a flat $GL(1) = GL(1,\mathbb{C})$ connection on $T_1\Sigma$ is uniquely parametrized by its $GL(1) \simeq \mathbb{C}^*$ -valued holonomies x_γ for $\gamma \in H_1(T_1\Sigma)$, with $x_{\gamma+\gamma'} = x_\gamma x_{\gamma'}$. We are requiring that the fiber holonomy is $x_u = -1$, whence the holonomies x_γ naturally become labelled by elements of twisted homology $\widetilde{H}_1(\Sigma)$.

Letting $\mathcal{X}_{\rm ab}[\Sigma]$ denote the space of standard (untwisted) $GL(1)$ flat connections on Σ, we note that there is an isomorphism $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma] \simeq \mathcal{X}_{\rm ab}[\Sigma]$ induced by any splitting $\widetilde{H}_1(\Sigma) \simeq H_1(\Sigma) \oplus \mathbb{Z}_2$. Such a splitting is given, for example, by the structure of a framed 3-manifold, with Σ the canonical cover of the boundary (Lemma 3). Equivalently, a choice of spin structure on Σ induces an isomorphism $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma] \simeq \mathcal{X}_{\rm ab}[\Sigma]$. Namely, a spin structure defines a (fiber-wise) 2-fold cover of $T_1\Sigma$; the pull-back of a twisted flat connection to the cover gives a connection with fiber holonomy $+1$, which may subsequently be pushed forward to Σ itself, providing the desired isomorphism (*cf.* [6, Sec. 10]).

Now suppose $\Sigma \xrightarrow{\pi} \mathcal{C}$ is the canonical double cover of the boundary of a framed 3-manifold. We want an odd version of $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma]$ whose coordinates are labelled by $\gamma \in \widetilde{H}_1^-(\Sigma)$. It can be defined as a projectivization, in the following sense. Let $\{\alpha_i\}_{i=1}^r$ be a basis for $H_1(\mathcal{C})$, and for an r-tuple of parameters $t = (t_1, .., t_r) \in (\mathbb{C}^*)^r$ and $\alpha = \sum_i a_i\alpha_i$ in $H_1(\mathcal{C})$ let $t^\alpha := t_1^{\alpha_1}\cdots t_r^{\alpha_r}$.

[19] Topological descriptions of non-abelianization were also discussed in [6, 37] and [40], which overlap with our constructions. In our language, [6, 37] considered moduli spaces of $PGL(K)$ (not just $PGL(2)$) connections on surfaces $\mathcal{C}$ consisting of an arbitrary big part $\mathcal{C}_{\rm big}$, with 2d ideal triangulation $\mathbf{t}_{\rm 2d}$, and all holes of $\mathcal{C}_{\rm big}$ filled by small discs. It was found for $PGL(2)$ that coordinates induced by non-abelianization coincide with Fock-Goncharov cluster coordinates. [40] generalized the $PGL(2)$ non-abelianization map to construct complex Fenchel-Nielsen coordinates for a surface $\mathcal{C}$ together with a pants decomposition. Such coordinates arise for us when $\mathcal{C}$ (as a framed boundary) contains big 3-holed spheres connected by small annuli, *cf.* [32]. Our definition of Φ in this case differs from that of [40], but is ultimately equivalent.

We set

$$\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma] := \widetilde{\mathcal{X}}_{\rm ab}[\Sigma]/(\mathbb{C}^*)^r\,, \qquad \text{with action } \; x_\gamma \mapsto t^{\pi_*\circ p(\gamma)} x_\gamma\,, \quad \gamma \in \widetilde{H}_1(\Sigma)\,,\; t \in (\mathbb{C}^*)^r\,. \tag{5.2}$$

Here $p(\gamma) \in H_1(\Sigma)$ is the projection to standard homology, and $\pi_* \circ p(\gamma) \in H_1(\mathcal{C})$ is the subsequent projection to the base. The coordinates invariant under the $(\mathbb{C}^*)^r$ action are precisely those x_γ with $\gamma \in \widetilde{H}_1^-(\Sigma)$. Thus the holonomies provide a map

$$x : \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma] \times \widetilde{H}_1^-(\Sigma) \to \mathbb{C}^*\,, \tag{5.3}$$

which is a homomorphism on the second factor, and non-degenerate in the sense that any basis $\{\gamma_i\}$ for $\widetilde{H}_1^-(\Sigma)$ induces an isomorphism

$$(x_{\gamma_i}) : \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma] \xrightarrow{\sim} (\mathbb{C}^*)^{\operatorname{rank} \widetilde{H}_1^-(\Sigma)} = (\mathbb{C}^*)^{\operatorname{rank} H_1^-(\Sigma)}\,. \tag{5.4}$$

Remark. One could also consider a space $\widetilde{\mathcal{X}}'_{\rm ab}[\Sigma]$ defined as the slice of $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma]$ on which $x_{P_+\gamma} = 1$ for all $\gamma \in \widetilde{H}_1(\Sigma)$. Our space $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ is a finite quotient of $\widetilde{\mathcal{X}}'_{\rm ab}[\Sigma]$ by $(\mathbb{Z}_2)^s$ (acting as multiplication by -1 on some $\mathbb{C}^*$ factors), where s is the number of $\mathbb{Z}_2$ factors in the torsion group $H_1(\Sigma)/(\operatorname{im} P_+ \oplus \ker P_+)$. Under the non-abelianization map of Section 5.3, $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ maps to framed flat $PGL(2)$ connections on a punctured base $\mathcal{C}^*$, whereas $\widetilde{\mathcal{X}}'_{\rm ab}[\Sigma]$ would map to twisted framed flat $SL(2)$ connections on $\mathbb{C}^*$.

The spaces $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma]$ and $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ have natural holomorphic Poisson and symplectic structures given by the Atiyah-Bott formula $\omega_{\rm ab} = \int_\Sigma \delta\mathscr{A} \wedge \delta\mathscr{A}$.[20] In this abelian setting, it is trivial to compute the Poisson brackets (on $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$, say)

$$\{\log x_\gamma, \log x_{\gamma'}\} = \langle \gamma, \gamma' \rangle\,, \qquad \gamma, \gamma' \in \widetilde{H}_1^-(\Sigma)\,. \tag{5.5}$$

To write the symplectic form in coordinates, we must choose some $\{\gamma_i\}_{i=1}^{\operatorname{rank} H_1^-(\Sigma)}$, $\gamma_i \in \widetilde{H}_1^-(\Sigma)$, lifting a basis $\{\overline{\gamma}_i\}$ of $H_1^-(\Sigma)$. Letting $\epsilon_{ij} = \langle \gamma_i, \gamma_j \rangle$ denote the non-degenerate intersection pairing and $x_i := x_{\gamma_i}$, we have

$$\omega_{\rm ab} = \frac{1}{2} \sum_{i,j} (\epsilon^{-1})^{ij} \frac{dx_i}{x_i} \wedge \frac{dx_j}{x_j}\,. \tag{5.6}$$

More elegantly, the symplectic form is induced by the Poincaré dual of the cup product (composed with a projection to odd homology)

$$\cup^* : \begin{array}{ccc} H_2(\Sigma) & \to & H_1^-(\Sigma) \wedge H_1^-(\Sigma) \\ [\Sigma] & \mapsto & \frac{1}{2}\sum_{ij} (\epsilon^{-1})^{ij} \overline{\gamma}_i \wedge \overline{\gamma}_j\,. \end{array} \tag{5.7}$$

From (5.7) we also obtain the canonical lift to K-theory (modulo torsion)

$$\hat{\omega}_{\rm ab} := \frac{1}{2} \sum_{i,j} (\epsilon^{-1})^{ij} x_i \wedge x_j \;\in\; K_2(\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma])_{\mathbb{Q}}\,. \tag{5.8}$$

[20] In this formula, $\mathscr{A}$ is a standard flat $GL(1)$ connection on Σ, represented locally as a 1-form. Implicitly, $\mathscr{A}$ is obtained by using *any* isomorphism between twisted and untwisted moduli spaces $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma] \simeq \mathcal{X}_{\rm ab}[\Sigma]$. The symplectic form is independent of the choice of isomorphism. This is manifest in expressions (5.5), (5.6), which are invariant under sign changes $x_\gamma \mapsto x_{\gamma+u} = -x_\gamma$.

5.2 Gluing abelian connections

Let M be any triangulated framed 3-manifold, and M' a framed 3-manifold glued by identifying faces of M. As usual we split the gluing into two parts $M \rightsquigarrow M_0 \rightsquigarrow M'$. The gluing gluing maps for spaces of flat connections $\widetilde{\mathcal{X}}^-_{\rm ab}$ on the canonical covers $\Sigma, \Sigma_0, \Sigma'$ of the respective boundaries. Due to the pairing (5.3), these gluing maps are automatically dual to the gluing maps in homology.

Concretely, first consider the gluing $M \rightsquigarrow M_0$, in which only interiors of faces of a big-boundary triangulation $\mathbf{t}_{2\rm d}$ are identified. The pre-images of each face in $T_1\Sigma$ retracts to a single S^1 fiber. Thus, any $GL(1)$ flat connection $\mathscr{A} \in \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ can be trivialized along the gluing region (aside from the universal fiber holonomy -1), and the gluing gives a map

$$\begin{array}{rccc} g^{(1)}_{GL(1)}: & \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma] & \twoheadrightarrow & \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma_0] \\ & x_{\widetilde{g}(\gamma)} & \mapsto & x_\gamma \qquad (\gamma \in \widetilde{H}^-_1(\Sigma_0))\,. \end{array} \tag{5.9}$$

The holonomies x_γ of $g^{(1)}_{GL(1)}(\mathscr{A})$ must be equal to the holonomies $x_{\widetilde{g}(\gamma)}$ of $\mathscr{A}$ itself, with $\widetilde{g}$ as in Section 2.5. Since $\widetilde{g}$ is an injection, $g^{(1)}_{GL(1)}$ is a surjection. (More precisely, there is an isomorphism $g^{(1)}_{GL(1)}: \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]/\operatorname{coker}\widetilde{g} \xrightarrow{\sim} \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma_0]$, with generators of the 2-torsion group $\operatorname{coker}\widetilde{g} \simeq \operatorname{coker} g$ acting as multiplication by -1 on $\mathbb{C}^*$ factors.) Moreover, since $\widetilde{g}$ preserves the intersection product, $g^{(1)}_{GL(1)}$ is a symplectomorphism.

Second, when filling in defects $M_0 \rightsquigarrow M'$, a flat connection on Σ_0 induces a flat connection on Σ' if and only if its holonomies along defect cycles $\nu \in \widetilde{G}'$ are trivial (again using the notation of Section 2.5). We denote this condition as "$x_{\nu\in\widetilde{G}'} = 1$." When it is satisfied, we may trivialize the flat connection in a neighborhood of the defects, then fill them in, obtaining a map

$$\begin{array}{rccc} g^{(2)}_{GL(1)}: & \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma_0]\big|(x_{\nu\in\widetilde{G}'}=1) & \twoheadrightarrow & \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma'] \\ & x_\gamma & \mapsto & x_{\widetilde{q}(\gamma)} \qquad (\gamma \in \widetilde{K}' \subset \widetilde{H}^-_1(\Sigma_0))\,. \end{array} \tag{5.10}$$

Now since $\widetilde{H}^-_1(\Sigma', \mathbb{Z}) = \widetilde{H}^-_1(\Sigma_0)//\widetilde{G}' = \widetilde{K}'/\widetilde{G}'$ (Prop. 1') we see that (5.10) is symplectic reduction, $\widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma'] \simeq \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma_0]//(\mathbb{C}^*)^{\operatorname{rank}\widetilde{G}'} = \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma_0]\big|_{x_{\nu\in\widetilde{G}'}=1}/(\mathbb{C}^*)^{\operatorname{rank}\widetilde{G}'}$, where the $(\mathbb{C}^*)^{\operatorname{rank}\widetilde{G}'}$ action is generated by using the $x_{\nu\in\widetilde{G}'}$ as moment maps (analogous to the action described in (4.32) following Theorem 1).

Combining the descriptions of (5.9)–(5.10), we deduce that the combined gluing map $g_{GL(1)} := g^{(2)}_{GL(1)} \circ g^{(1)}_{GL(1)}$ is a symplectic reduction of a finite quotient, governed by gluing equations

$$x_{\widetilde{g}(\gamma)} = x_{\widetilde{q}(\gamma)}\,, \qquad \forall\, \gamma \in \widetilde{K}'\,. \tag{5.11}$$

When M and M' are defect-free, we may use Lemma 5 to replace $\widetilde{G}', \widetilde{K}'$ by the finite-index subgroups $\widetilde{G} := \widetilde{h}(\mathbf{P}_G)$ and $\widetilde{K} := \widetilde{h}(\mathbf{P}_0 \oplus \mathbb{Z}_2)$ corresponding to the path algebra on $(\mathcal{C}_0)_{\rm small}$. Then

$$g_{GL(1)}: \mathcal{X}^-_{\rm ab}[\Sigma]\big|(x_{\widetilde{g}(\widetilde{G})} \equiv 1) \twoheadrightarrow \mathcal{X}_{\rm ab}[\Sigma'] \simeq (\mathcal{X}^-_{\rm ab}[\Sigma]/Z)//(\mathbb{C}^*)^{\operatorname{rank}\widetilde{G}}\,, \tag{5.12}$$

governed by gluing equations identical to (4.29): $x_{\tilde{g}(\gamma)} = x_{\tilde{q}(\gamma)}$ for all $\gamma \in \widetilde{K}$. (The torsion group $Z \simeq \widetilde{H}_1^-(\Sigma)/\widetilde{H}$ acts on $\mathcal{X}_{\rm ab}^-[\Sigma]$ the same way as in Theorem 1.)

5.3 Spectral networks and non-abelianization

Suppose that $\mathcal{C}$ is the boundary of a framed 3-manifold without defects, with canonical cover $\Sigma \xrightarrow{\pi} \mathcal{C}$, and that $\pi_1(\mathcal{C}_{\rm small})$ is abelian. Fix a triangulation $\mathfrak{t}_{\rm 2d}$ of $\mathcal{C}_{\rm big}$. By the results of Section 5.1, the symplectic moduli space of twisted abelian flat connections on Σ, $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma] \simeq (\mathbb{C}^*)^{\operatorname{rank} H_1^-(\Sigma)}$ is coordinatized by x_γ, $\gamma \in \widetilde{H}_1^-(\Sigma)$. Define the open subset

$$\mathcal{P}_{\rm ab}[\Sigma] := \left\{ \mathscr{A} \in \widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma] \,\middle|\, x_\lambda^{(a)}(\mathscr{A}) \neq 1,\ (x_\alpha^{(t)}(\mathscr{A}), x_\beta^{(t)}(\mathscr{A})) \neq (1,1) \right\} \subset \widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma]\,, \tag{5.13}$$

on which holonomies around tori and annuli are never totally unipotent. This is analogous to the definition of the $PGL(2)$ moduli space $\mathcal{P}[\mathcal{C}, \mathfrak{t}] \subset \mathcal{X}[\mathcal{C}]$ on p. 33. In this section we construct a non-abelianization map

$$\Phi[\mathfrak{t}_{\rm 2d}] :\ \mathcal{P}_{\rm ab}[\Sigma] \to \mathcal{P}[\mathcal{C}; \mathfrak{t}_{\rm 2d}]\,, \tag{5.14}$$

defined via a topological "spectral network" that (slightly) extends that of [6] (*cf.* Footnote 19). We show that

Proposition 4 *With notation as above, the map $\Phi[\mathfrak{t}_{\rm 2d}]$ is 1-1 and a (holomorphic) K_2 symplectomorphism. It identifies twisted $GL(1)$ holonomies x_γ with the path-functions x_γ on $\mathcal{P}[\mathcal{C}; \mathfrak{t}_{\rm 2d}]$ labelled by elements $\gamma \in \widetilde{H}_1^-(\Sigma)$; in other words $\Phi[\mathfrak{t}_{\rm 2d}]^*(x_\gamma) = x_\gamma$.*

We construct Φ in two steps. First, given a twisted abelian flat connection $\mathscr{A}$ on Σ, we can push it forward to a framed flat $PGL(2)$ connection $\pi_*\mathscr{A}$ on $\mathcal{C}\backslash\mathfrak{b}$, *i.e.* in the complement of the branch locus,

$$\pi_* : \widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma] \to \mathcal{X}[\mathcal{C}\backslash\mathfrak{b}]\,. \tag{5.15}$$

To see this, note that the projection $\pi : \Sigma \to \mathcal{C}$ can be extended uniquely to a bundle map $\pi : T_1(\Sigma\backslash\mathfrak{b}) \to T_1(\mathcal{C}\backslash\mathfrak{b})$ that is globally two-to-one and an isomorphism on the unit-tangent fibers. Then a flat line bundle $L \to T_1\Sigma$ induces a flat rank-two bundle $E' \to T_1(\mathcal{C}\backslash\mathfrak{b})$, locally of the form $L^+ \oplus L^-$, where $L^\pm$ denote (locally) the bundles over the two sheets of the cover $\pi : T_1(\Sigma\backslash\mathfrak{b}) \to T_1(\mathcal{C}\backslash\mathfrak{b})$. The induced flat connection $\mathscr{A}'$ on E' (locally of the form $\mathscr{A}^+ \oplus \mathscr{A}^-$) has holonomy valued in $GL(2)$, and equal to $-\mathbb{1}$ around unit-tangent fibers. Taking its projectivization, we obtain a connection $\mathscr{A}''$ of $PGL(2)$ holonomy. Since $-\mathbb{1} \simeq \mathbb{1}$ in $PGL(2)$, the holonomy of $\mathscr{A}''$ around unit-tangent fibers is trivial, so $\mathscr{A}''$ descends to a flat $PGL(2)$ connection on a rank-two bundle $E_* \to (\mathcal{C}\backslash\mathfrak{b})$, which we call $\pi_*\mathscr{A}$.

A priori, this procedure describes a map π_* from $\widetilde{\mathcal{X}}_{\rm ab}[\Sigma]$ to $\mathcal{X}[\mathcal{C}\backslash\mathfrak{b}]$. Our definition of the "odd" space $\widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma]$ in (5.2), however, was precisely engineered so that the map π_* would factor through to $\widetilde{\mathcal{X}}_{\rm ab}^-[\Sigma]$. Indeed, the action $x_\gamma \mapsto t^{\pi_* \circ p(\gamma)} x_\gamma$ in (5.2) is just a lift of the projectivization action on E'.

The connection $\pi_*\mathscr{A}$ is also naturally framed. Recall from Section 2.4 that the canonical cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ can be glued together from two sheets Σ^+ and Σ^-, each a copy of $\mathcal{C}\backslash\Gamma_{\rm br}$. We correlate the local decomposition $E' = L^+ \oplus L^-$ above with the labeling of these sheets. Then induces a canonical diagonal decomposition $E_* = L^+ \oplus L^-$ (as a projectively flat bundle) over each connected component of $\mathcal{C}\backslash\Gamma_{\rm br}$, and in particular over $\mathcal{C}_{\rm small}$. We take the flag $L^+ \subset E_*$ to be the framing data on each component of $\mathcal{C}_{\rm small}$.

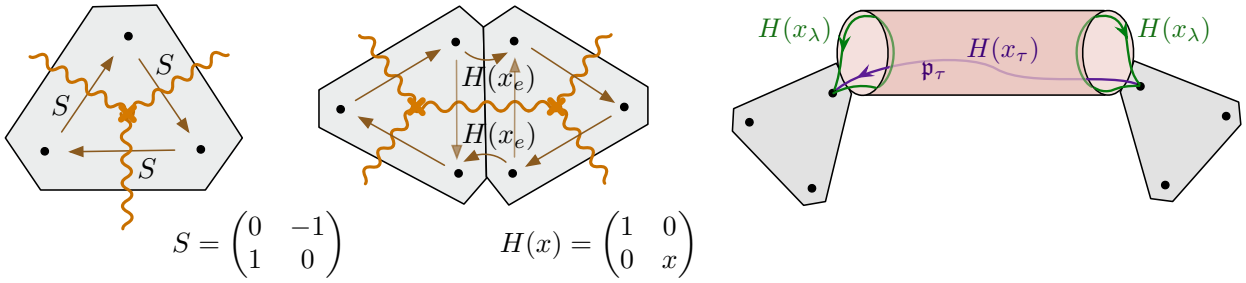

Figure 26. Local frames and parallel transport for the flat connection $\pi_*\mathscr{A}$

The second step is to modify the push-forward connection $\pi_*\mathscr{A}$, defined over $\mathcal{C}\backslash\mathfrak{b}$, to obtain a new (and more interesting) connection that extends over the branch points but may have unipotent singularities at punctures on the small boundary. We proceed as follows.

For small-sphere and small-torus components of $\mathcal{C}$, no modification is required. The remaining connected components of $\mathcal{C}$ consist of triangulated big boundary with holes filled in by small discs and annuli. We then choose local frames for the projectively flat bundle $E_* \to \mathcal{C}\backslash\mathfrak{b}$ over three points p in each face t of $\mathfrak{t}_{\rm 2d}$ so that the parallel transport of $\pi_*\mathscr{A}$ is given by the transformations in Figure 26. Locally, each frame is a choice of vectors in the lines $L^\pm$. Thus the parallel transport is diagonal over connected components of $\mathcal{C}\backslash\Gamma_{\rm br}$, given by matrices $H(x_\mathfrak{p}) = \mathrm{diag}\,(1, x_\mathfrak{p})$ for appropriate functions $x_\mathfrak{p} = (x_e, x_\tau, x_\lambda)$ on $\widetilde{X}^-_{\rm ab}[\Sigma]$. Across branch cuts, the lines are exchanged and the parallel transport $S = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ is anti-diagonal.[21] The nontrivial holonomy $S^3 = S$ around branch points is what prevents E_* from extending over the branch points.

To fix this we introduce a (topological) *spectral network* $\mathcal{W} = \mathcal{W}[\mathfrak{t}_{\rm 2d}, \{\mathfrak{p}_\tau\}]$ on $\mathcal{C}$. It depends on $\mathfrak{t}_{\rm 2d}$ and has a mild dependence on a choice of paths $\mathfrak{p}_\tau$ traversing small annuli. Roughly, $\mathcal{W}$ is an unoriented graph dual to small edges of $\mathfrak{t}_{\rm 2d}$ and to the paths $\mathfrak{p}_\tau$, as in Figure 27. Formally, $\mathcal{W}$ has vertices at all branch points and at one point in the interior of each small disc and annulus. Its edges connect the vertex in each small disc d to every branch

[21] A similar choice of local frames was used in the non-abelianization constructions of [6, 37]. Note that the sign in $S = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ is particularly natural given the push-forward construction of $\pi_*\mathscr{A}$. Namely, since $\mathscr{A}$ has holonomy -1 around unit-tangent fibers of $T_1\Sigma$, it follows that the holonomy $h_{\rm br}$ of $\mathscr{A}'$ around any cycle surrounding a branch point (and wrapping the unit-tangent fiber of $T_1\mathcal{C}$ any number of times) must satisfy $h_{\rm br}^2 = -\mathbf{1} \in GL(2)$. To be compatible with the local decomposition $E' = L^+ \oplus L^-$, it must also be anti-diagonal, and can be chosen (modulo gauge equivalence) as $h_{\rm br} = \pm\left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$. After projectivizing, we obtain holonomy $S^3 = S = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right) \in PGL(2)$ around branch points.

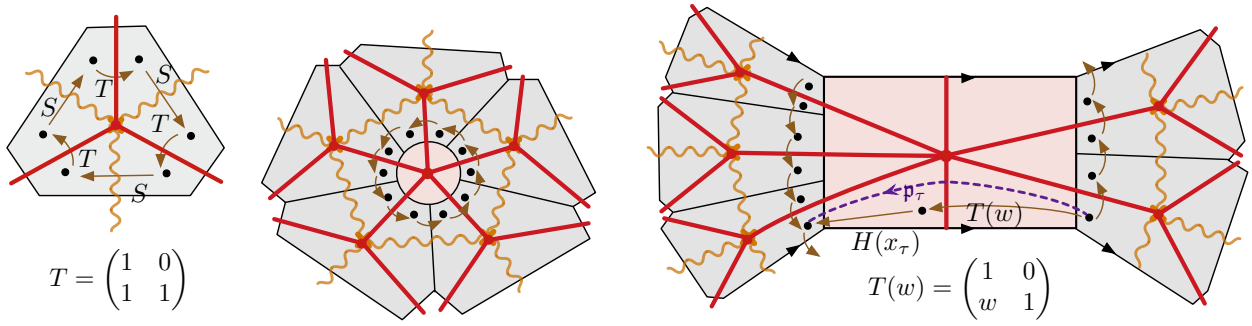

Figure 27. Structure of the spectral network $\mathcal{W}$ (in red) on a face of $\mathbf{t}_{2d}$ (left), in the neighborhood of a small disc (center), and in the neighborhood of a small annulus (right). We indicate the unipotent modifications of $\pi_* \mathscr{A}$ corresponding to the edges of the spectral network on faces (left) and wrapping around annuli (right).

point in a face of $\mathbf{t}_{2d}$ surrounding d; connect the vertex in each small annulus a to the branch points in the faces surrounding a, without crossing $\mathfrak{p}_\tau^{(a)}$; and connect the vertex in a to itself via a circular path in the homology class of $\mathfrak{p}_\lambda^{(a)}$.

We perform unipotent modifications of E_* across walls of the spectral network, in two rounds. First, note that each edge s of $\mathcal{W}$ ending at a branch point $b \in \mathfrak{b}$ passes close to a single point p labeling a frame for E_*. We add a unipotent modification by $T = \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$ when moving across s clockwise from the perspective of b — meaning explicitly that we split the one frame over p into two, with new parallel transport T between them. This creates a new projectively flat bundle E'_*. Its holonomy around any branch point is trivial by virtue of the relation $(ST)^3 = \mathbb{1}$, hence E'_* extends over the branch locus. However, E'_* may have nontrivial unipotent holonomy at the vertices of $\mathcal{W}$ on small discs and annuli.

Now, for each small annulus a, we add a second unipotent modification to E'_* by $T(w_a) := \left(\begin{smallmatrix} 1 & 0 \\ w_a & 1 \end{smallmatrix}\right)$ across the edge of $\mathcal{W}$ wrapped around a (homologous to $\mathfrak{p}_\lambda^{(a)}$). For concreteness, we take this modification to split either of the two frames of E'_* at the tail of $\mathfrak{p}_\tau^{(a)}$. (These frames were already split by the first modification; since $[H(w_a), H(1)] = 0$, it does not matter which we take.) Then, as long as $x_\lambda^{(a)} \neq 1$, there is a unique $w_a \in \mathbb{C}$ that trivializes the holonomy around the vertex of $\mathcal{W}$ on a. (The calculation is identical to solving for t in (B.3), p. 62.)

Following these two rounds of unipotent modifications, we arrive at a projectively flat bundle E that extends over all of $\mathcal{C}$ except the vertices of $\mathcal{W}$ on small discs — *i.e.* over the "punctured boundary" $\mathcal{C}^*$ of Section 4.1. Let $\Phi(\mathscr{A})$ denote the flat $PGL(2)$ connection on E. Since unipotent modifications preserve the flag $L^+ \subset E$ over each component of C_{small}, $\Phi(\mathscr{A})$ is a naturally a framed flat connection. Thus, we've defined a map

$$\Phi : \widetilde{\mathcal{X}}^-_{\text{ab}}[\Sigma]\big|_{x_\lambda \neq 1} \to \mathcal{X}[\mathcal{C}]\,. \tag{5.16}$$

Its definition depends both on a triangulation $\mathbf{t}_{2d}$ and on a choice of paths $\mathfrak{p}_\tau$. We will prove momentarily that the dependence on $\mathfrak{p}_\tau$'s is trivial and that when Φ is restricted to both $x_\lambda \neq 1$ (for annuli) and $(x_\alpha, x_\beta) \neq (1, 1)$ (for tori) its image is precisely $\mathcal{P}[\mathcal{C}, \mathbf{t}_{2d}]$.

Proof of Prop. 4 Given $\mathscr{A} \in \widetilde{\mathcal{X}}^-_{\rm ab}[\Sigma]$ with $x_\lambda \neq 1$ and $(x_\alpha, x_\beta) \neq 1$, observe that $\Phi(\mathscr{A})$ is a framed flat $PGL(2)$ connection on $\mathcal{C}^*$ with exactly the same set of local frames and parallel transports as are used in Appendix B to uniquely reconstruct a connection $\mathcal{A} \in \mathcal{P}[\mathcal{C}, \mathbf{t}_{\rm 2d}]$ from its coordinates. Thus $\Phi(\mathscr{A}) \in \mathcal{P}[\mathcal{C}, \mathbf{t}_{\rm 2d}]$. Indeed, the path-coordinates $x_{\mathfrak{p}}$ of $\Phi(\mathscr{A}) \in \mathcal{P}[\mathcal{C}]$, showing up as diagonal entries of parallel transport matrices, are manifestly equal to the corresponding path-coordiantes $x_{\mathfrak{p}}$ labeling the abelian holonomies of $\mathscr{A}$. (In each case, $x_{\mathfrak{p}}$ depends only on the twisted homology class $\widetilde{h}(\mathfrak{p}) \in \widetilde{H}_1^-(\Sigma)$.) From (5.4) and Prop. 2 it follows that $\Phi : \mathcal{P}_{\rm ab}[\Sigma] \to \mathcal{P}[\mathcal{C}, \mathbf{t}_{\rm 2d}]$ is one-to-one.

Since path-functions $x_{\mathfrak{p}}$ uniquely determine $\Phi(\mathscr{A})$, independent of any particular choice of traversing paths $\mathfrak{p}_\tau$ for the annuli, the map Φ cannot depend on the choice of $\mathfrak{p}_\tau$'s used in unipotent modification. Thus it depends at most on a big-boundary triangulation $\mathbf{t}_{\rm 2d}$.

The fact that Φ is a K_2 symplectomorphism follows immediately by comparing the K_2 forms (4.17), (5.8) (*cf.* the Poisson brackets (4.18), (5.5)), which look identical in x_γ coordiantes, and are both controlled by the intersection form on $H_1^-(\Sigma)$. There is also an alternative, coordinate-free, proof of the fact that Φ is a symplectomorphism. Namely, we decompose Φ as a composition of push-forward (π_*) and unipotent modifications. The push-forward map is obviously a symplectomorphism for the Atiyah-Bott symplectic/Poisson structures. It was then shown in [6, Section 10.4] that unipotent modification preserves the holomorphic symplectic structure. □

5.4 Non-abelianization commutes with gluing

We combine the results of the previous sections in a final theorem about gluing. Suppose that we glue framed 3-manifolds $M \rightsquigarrow M_0 \rightsquigarrow M'$ as in Section 2.5, where M and M' have no defects, and their small boundaries $\mathcal{C}_{\rm small}$ and $\mathcal{C}'_{\rm small}$ have only discs, annuli, and tori. Let us fix compatible big-boundary triangulations $\mathbf{t}_{\rm 2d}$ for $\mathcal{C}_{\rm big}$ and $\mathbf{t}'_{\rm 2d}$ for $(\mathcal{C}_0)_{\rm big} = \mathcal{C}'_{\rm big}$.

Theorem 2 *Gluing and non-abelianization maps fit into a commutative diagram*

$$\begin{array}{ccc} \mathcal{P}_{\rm ab}[\Sigma]\big|_{x_{\tilde{g}(\tilde{G})} \equiv 1,\, R'} & \overset{g_{GL(1)}}{\twoheadrightarrow} & \mathcal{P}_{\rm ab}[\Sigma'] \;\; = (\mathcal{P}_{\rm ab}[\Sigma]_{R'}/Z) /\!\!/ (\mathbb{C}^*)^{\operatorname{rank}\tilde{G}} \\ \Phi[\mathbf{t}_{\rm 2d}] \downarrow & & \Phi[\mathbf{t}'_{\rm 2d}] \downarrow \\ \mathcal{P}[\mathcal{C}; \mathbf{t}_{\rm 2d}]\big|_{x_{\tilde{g}(\tilde{G})} \equiv 1,\, R'} & \overset{g_{PGL(2)}}{\twoheadrightarrow} & \mathcal{P}[\mathcal{C}'; \mathbf{t}'_{\rm 2d}] = (\mathcal{P}[\mathcal{C}; \mathbf{t}_{\rm 2d}]_{R'}/Z) /\!\!/ (\mathbb{C}^*)^{\operatorname{rank}\tilde{G}}, \end{array} \tag{5.17}$$

where R' is the technical restriction $x_{\tilde{g}\tilde{q}^{-1}(\lambda)} \neq 1$, $(x_{\tilde{g}\tilde{q}^{-1}(\alpha)}, x_{\tilde{g}\tilde{q}^{-1}(\beta)}) \neq (1,1)$ and $Z \simeq \widetilde{H}_1^-(\Sigma)/\tilde{H}$ (as in Theorem 1). The vertical maps are 1-1 K_2 symplectomorphisms and the horizontal maps are K_2 symplectic reduction of finite quotients. Thus the trivial symplectic reduction in the gluing of $GL(1)$ moduli spaces on Σ (5.12) induces the non-trivial symplectic reduction (Theorem 1) in the gluing of $PGL(2)$ moduli spaces on $\mathcal{C}$. Both reductions are governed by the same gluing equations

$$x_{\tilde{g}(\gamma)} \;=\; x_{\tilde{q}(\gamma)}\,, \qquad \forall\, \gamma \in \tilde{K} \subset \tilde{H}_1^-(\Sigma_0)\,, \tag{5.18}$$

ultimately arising from the isomorphism $\widetilde{H}_1^-(\Sigma') = \widetilde{H}//\widetilde{g}(\widetilde{G}) = \widetilde{g}(\widetilde{K})/\widetilde{g}(\widetilde{G})$ *with* $\widetilde{H} \subset \widetilde{H}_1^-(\Sigma)$ *of finite index (Lemma 5, Prop. 1').*

Proof. Commutativity of the square (5.17) follows easily by comparing the description of abelian gluing in Section 5.2 to the description of $PGL(2)$ gluing in Section 4.4. Alternatively (and more explicitly), we have already seen in (4.29) and (5.12) that both the $GL(1)$ and $PGL(2)$ gluing maps are governed by the same gluing equations (5.18), labelled by elements of twisted odd homology; since the non-abelianization maps Φ preserve the x_γ functions (by Prop. 4), it follows that the square must commute. The remaining claims follow immediately from the result of Prop. 4 that the non-abelianization maps are 1-1 symplectomorphisms. □

Acknowledgements

We especially wish to thank D. Futer, L. Hollands, A.B. Goncharov, G. Moore, A. Neitzke, and C. Zickert for helpful conversations and advice. We are also indebted to Center for Quantum Geometry of Moduli Spaces for their hospitality during the workshop *Pressure Metrics and Higgs Bundles*, Aug 19-22, 2013 (partial funding through the GEAR network), where the ideas for this project were conceived. T.D. is supported by DOE grant DE-SC0009988, and in part by the European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013), ERC Grant agreement 335739. R.V. is supported by a VENI grant from the Netherlands Organization for Scientific Research.

A Odd results

In this section, we collect some basic results about odd homology, and review their proofs. Notation is as in Section 2.1. In particular, $\Sigma \xrightarrow{\pi} \mathcal{C}$ is an oriented double cover of a closed, oriented surface, branched over a locus of points $\mathfrak{b}$.

A.1 Basics

Let $\sigma : \Sigma \to \Sigma$ denote the deck-transformation homeomorphism. It preserves orientation and its fixed-point locus is precisely $\mathfrak{b}(\pi)$. It induces a push-forward automorphism on homology groups $\sigma_* : H_\bullet(\Sigma) \xrightarrow{\sim} H_\bullet(\Sigma)$. Letting $P_\pm := \mathbb{1} \pm \sigma_*$, we define

$$H_\bullet^-(\Sigma) := \ker P_+ \,, \qquad H_\bullet^+(\Sigma) := \ker P_- \,. \tag{A.1}$$

Notice that

$$P_\pm^2 = 2P_\pm \,, \qquad P_\pm P_\mp = 0 \,, \qquad P_+ + P_- = 2\,\mathbb{1} \,. \tag{A.2}$$

Thus, $P_\pm$ are close to being orthogonal projection operators. They fail to be proper projections due to the factors of 2 in (A.2).

Lemma 6 *The quasi-projections $P_\pm$ obey the following properties:*

a) $\operatorname{im} P_\pm \subset \ker P_\mp$*;*

b) $\ker P_+ \cap \ker P_- = \{0\}$*;*

c) $\ker P_\mp / \operatorname{im} P_\pm \simeq (\mathbb{Z}_2)^{r_\pm}$ *(for some $r_\pm$) is finite, containing only 2-torsion;*

d) $H_\bullet / (\ker P_+ \oplus \ker P_-)$ *is also finite, containing only 2-torsion.*

Proof. Parts (a-b) follows trivially from (A.2). Parts (c-d) follow because $P_\pm$ can be promoted to honest orthogonal projections after tensoring with $\mathbb{Q}$. More concretely, given any $\alpha \in H_\bullet(\Sigma)$ such that $\alpha = 2\beta$ for some $\beta \in H_\bullet(\Sigma)$ we can uniquely decompose $\alpha = \alpha_+ + \alpha_-$ with $\alpha_\pm \in \operatorname{im} P_\pm$ (namely, $\alpha_\pm = P_\pm \beta$), proving (d); and given any $\beta \in \ker P_\mp$ we have $2\beta = P_\pm \beta$, proving (c). □

The first homology group $H_1(\Sigma)$ has a non-degenerate skew-symmetric intersection form $\langle\ ,\ \rangle : \bigwedge^2 H_1(\Sigma, \mathbb{Z}) \to \mathbb{Z}$, which is preserved by σ_* (since σ is an orientation-preserving homeomorphsim). Thus it is also preserved by $P_\pm$. We find that

Lemma 7 *Letting $\ker P_\pm, \operatorname{im} P_\pm$ denote kernels and images in $H_1(\Sigma)$ now, we have*

a) $\langle \ker P_\pm, \ker P_\mp \rangle = 0$ *;*

b) the intersection form is non-degenerate on $\ker P_+$ *and on* $\ker P_-$ *;*

c) the intersection form is even on $\operatorname{im} P_+$ *and* $\operatorname{im} P_-$ *;*

d) if $\langle \alpha, \beta \rangle = 0$ *for all* $\beta \in \operatorname{im} P_+$ *then* $\alpha \in \ker P_+$ *(and similarly with* $+ \leftrightarrow -$*).*

Proof. These are all simple consequences of invariance of the intersection form under $\sigma_*, P_\pm$. For example, if $\alpha_\pm \in \ker P_\pm$ then $\langle \alpha_+, \alpha_- \rangle = \langle \sigma\alpha_+, \sigma\alpha_- \rangle = -\langle \alpha_+, \alpha_- \rangle$, hence $\langle \alpha_+, \alpha_- \rangle = 0$, proving (a). Nondegeneracy on $H_1(\Sigma)$ then implies (b) and (d). For (c) note that if $\alpha, \beta \in \operatorname{im} P_+$ (say) then $\langle \alpha, \beta \rangle = \langle P_+\gamma, \beta \rangle = \langle \gamma, P_+\beta \rangle = 2\langle \gamma, \beta \rangle \in 2\,\mathbb{Z}$. □

Let $\pi_* : H_1(\Sigma) \to H_1(\mathcal{C})$ denote the induced action of the projection $\Sigma \xrightarrow{\pi} \mathcal{C}$. For any oriented curve $\gamma \subset (\mathcal{C} \backslash \mathfrak{b})$ (in the complement of the branching locus), the preimage $\pi^{-1}(\gamma)$ consists of one or two oriented curves on Σ. Let $\ell^+(\gamma)$ denote the homology class $[\pi^{-1}(\gamma)] \in H_1(\Sigma)$, and notice that it only depends on the homology class $[\gamma] \in H_1(\mathcal{C})$. (It is clear that $\ell^+(\gamma)$ is invariant under local homotopies of γ that do not cross the branching locus $\mathfrak{b}$. If a homotopy crosses $\mathfrak{b}$, invariance of $\ell^+(\gamma)$ is illustrated in Figure 28.) Thus the "even lift" ℓ^+ can be extended by linearity to a map $\ell^+ : H_1(\mathcal{C}) \to H_1(\Sigma)$. Moreover, we have

$$\pi_* \circ \ell^+ = 2\mathbb{1}\,, \qquad \ell^+ \circ \pi_* = P_+\,, \qquad P_- \circ \ell^+ = 0\,. \tag{A.3}$$

Lemma 8 *a)* $H_1^-(\Sigma) = \ker \pi_*|_{H_1(\Sigma)}$ *;*

b) $\ell^+ : H_1(\mathcal{C}) \hookrightarrow H_1^+(\Sigma)$ *is an injection with finite (2-torsion) cokernel;*

c) thus $H_1(\Sigma, \mathbb{Q}) = H_1^+(\Sigma, \mathbb{Q}) \oplus H_1^-(\Sigma, \mathbb{Q}) \simeq H_1(\mathcal{C}, \mathbb{Q}) \oplus H_1^-(\Sigma, \mathbb{Q})$*, and in particular*

$$\operatorname{rank} H_1^-(\Sigma) = \operatorname{rank} H_1(\Sigma) - \operatorname{rank} H_1(\mathcal{C})\,.$$

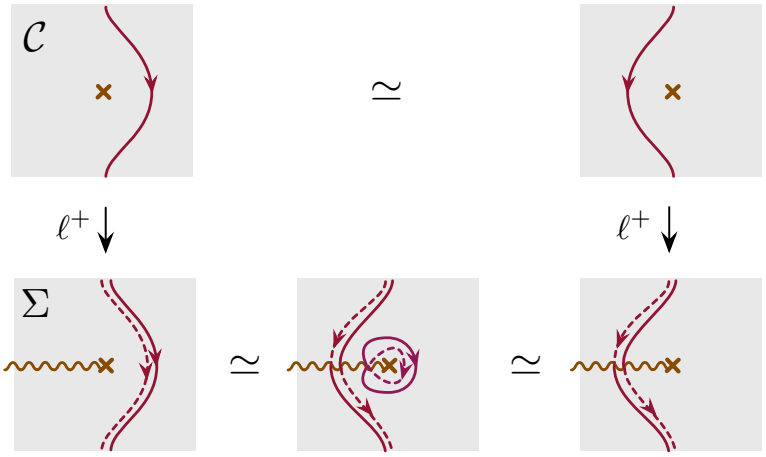

Figure 28. Demonstrating the invariance of the even lift $\ell^+(\gamma) = [\pi^{-1}(\gamma)]$ under a homotopy of γ through a branch point on $\mathcal{C}$. We represent Σ locally by drawing a branch cut emanating from the branch point, and draw the curves $\pi^{-1}(\gamma) \subset \Sigma$ using solid and dashed lines for the two sheets of the cover. On the bottom row of the figure, "$\simeq$" denotes equivalences in $H_1(\Sigma)$.

Proof. From $\pi_* \circ \ell^+ = 2\mathbb{1}$ we see that ℓ^+ is injective, so from $\ell^+ \circ \pi_* = P_+$ it follows that $\ker \pi_*|_{H_1(\Sigma)} = \ker P_+|_{H_1(\Sigma)}$, proving (a). From $P_- \circ \ell^+ = 0$ (or simply from the definition of ℓ^+) we see that $\operatorname{im} \ell_+ \subset H_1^+(\Sigma)$. Moreover, given any $\alpha \in H_1^+(\Sigma)$ we have $2\alpha = P_+\alpha = \ell^+(\pi_*\alpha)$, hence $2\alpha \in \operatorname{im} \ell_+$, proving (b). Part (c) follows from (b) and Lemma 6d. $\square$

A.2 Chain complexes

To analyze the effect of gluing on odd homology, it is convenient to have odd versions of standard exact sequences. The following result shows that we can restrict exact sequences to odd homology, modulo 2-torsion.

Lemma 9 *Let* $(A_\bullet, \delta_\bullet) = \big[\to A_i \xrightarrow{\delta_i} A_{i-1} \xrightarrow{\delta_{i-1}} A_{i-2} \to \big]$ *be a chain complex of abelian groups. Let* σ *be an involution of* $A_\bullet$ *that preserves grading and commutes with* $\delta_\bullet$*. Set* $P_+ := \mathbb{1} + \sigma$ *and* $A_\bullet^- := \ker P_+|_{A_\bullet}$*, and let* $\delta_\bullet^-$ *be the restriction of* $\delta_\bullet$ *to* $A_\bullet^-$*. Then*

a) $(A_\bullet^-, \delta_\bullet^-)$ *is also a chain complex;*

b) σ *induces an involution on homology* $H_\bullet(A_\bullet, \delta_\bullet)$*, and letting* $H_\bullet^-(A_\bullet, \delta_\bullet) := \ker P_+|_{H_\bullet(A_\bullet, \delta_\bullet)}$*, there is an isomorphism* $H_\bullet(A_\bullet^-, \delta_\bullet^-) \simeq H_\bullet^-(A_\bullet, \delta_\bullet)$ *modulo 2-torsion;*

c) If $(A_\bullet, \delta_\bullet)$ *is exact and injective on the left, say* $0 \to A_d \xrightarrow{\delta_d} A_{d-1} \xrightarrow{\delta_{d-1}} A_{d-2} \to \ldots$ *for some* d*, then* $(A_\bullet^-, \delta_\bullet^-)$ *is exact in the first two places (at* A_d^- *and* A_{d-1}^-*) and its homology is 2-torsion thereafter.*

Proof. For (a), observe that if $\alpha \in A_i^-$ then $P_+(\delta_i \alpha) = \delta_i(P_+\alpha) = 0$. Therefore, $\operatorname{im} \delta_i^- \subset \operatorname{im} \delta_i \cap A_{i-1}^- \subset \ker \delta_{i-1}^-$, showing that $(A_\bullet^-, \delta_\bullet^-)$ is a complex.

Next, observing that σ fixes both $\ker \delta_\bullet$ and $\operatorname{im} \delta_\bullet$, since it commutes with $\delta_\bullet$. Thus there is an induced involution on classes in the i-th homology group $H_i = \ker \delta_i / \operatorname{im} \delta_{i+1}$, given by

$\sigma[\alpha] := [\sigma\alpha]$, and it makes sense to consider

$$H_i^- := \ker P_+|_{H_i} = \langle\, \alpha \in \ker \delta_i \,|\, P_+\alpha \in \operatorname{im} \delta_{i+1} \,\rangle / \operatorname{im} \delta_{i+1}\,. \tag{A.4}$$

Any class $[\alpha] \in H_i^-$ can be represented by $\alpha \in \ker \delta_i$ with $P_+\alpha = \delta_{i+1}\beta$. We apply P_+ to this relation to find $2P_+\alpha = P_+(\delta_{i+1}\beta)$. Letting $\alpha' = 2\alpha - \delta_{i+1}\beta$ we see that $2[\alpha] \simeq [\alpha']$ and $P_+\alpha' = 0$. Thus two times every class in H_i^- has a representative annihilated by P_+. This means that modulo 2-torsion

$$H_i^- \simeq (\ker \delta_i \cap A_i^-)/(\operatorname{im} \delta_{i+1} \cap A_i^-) = \ker \delta_i^- / (\operatorname{im} \delta_{i+1} \cap A_i^-)\,. \tag{A.5}$$

(More precisely, the RHS injects into the LHS with finite cokernel.) Moreover, if $\alpha \in \operatorname{im} \delta_{i+1} \cap A_i^-$, *i.e.* $\alpha = \delta_{i+1}\beta$ and $P_+\alpha = 0$, then $2\alpha = P_-\alpha = \delta_{i+1}(P_-\beta) \in \operatorname{im} \delta_{i+1}^-$. Therefore, the quotient $(\operatorname{im} \delta_{i+1} \cap A_i^-)/\operatorname{im} \delta_{i+1}^-$ is 2-torsion, and

$$H_i^- \simeq \ker \delta_i^- / \operatorname{im} \delta_{i+1}^- = H_i(A_\bullet^-, \delta_\bullet^-) \tag{A.6}$$

modulo 2-torsion, as needed for part (b).

Finally, suppose $(A_\bullet, \delta_\bullet)$ is an exact sequence as in (c). Part (b) implies that $H_\bullet(A_\bullet^-, \delta_\bullet^-)$ is 2-torsion; but we can do better. Since δ_d is injective, δ_d^- is also injective, so $H_d(A_\bullet^-, \delta_\bullet^-) = 0$. Moreover, $\ker \delta_{d-1}^- = \ker \delta_{d-1} \cap A_i^- = \operatorname{im} \delta_d \cap A_i^- = \operatorname{im} \delta_d^-$ (the last equality again follows because δ_d is injective: if $\alpha = \delta_d\beta$ and $P_+\alpha = 0$ then $\delta_d(P_+\beta) = 0$, hence $P_+\beta = 0$), so $H_{d-1}(A_\bullet^-, \delta_\bullet^-) = 0$. □

A.3 Cellular description

For a branched cover $\Sigma \xrightarrow{\pi} \mathcal{C}$, we can describe a set of generators for $H_1^-(\Sigma)$ and the intersection form on them very explicitly. We do so by applying Lemma 9 to a cell complex for Σ.

Choose a finite cell decomposition of $\mathcal{C}$ with 2-cells $\{f_i\}$, 1-cells $\{e_i\}$ and 0-cells $\{p_i\}$, such that every branch point of the cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ is a 0-cell. By lifting to the two sheets of Σ, this induces a cell decomposition of Σ with 2-cells $D_2 = \{f_i^+, f_i^-\}$, 1-cells $D_1 = \{e_i^+, e_i^-\}$, and 0-cells $D_0 = \{p_i^+, p_i^- \,|\, p_i \notin \mathfrak{b}\} \cup \mathfrak{b}$ (where superscripts $\pm$ indicate local choices of lifts). We immediately recover from this the standard Riemann-Hurwitz formula

$$\chi(\Sigma) = 2\chi(\mathcal{C}) - \#(\mathfrak{b})\,, \tag{A.7}$$

and in combination with Lemma 8(c) we get

Lemma 10 *Suppose that $\mathcal{C}$ and Σ are both connected. Then their genera are related by $g(\Sigma) = 2g(\mathcal{C}) + \#(\mathfrak{b})/2 - 1$ and*

$$\begin{aligned} \operatorname{rank} H_1(\Sigma) &= 2 \operatorname{rank} H_1(\mathcal{C}) + \#(\mathfrak{b}) - 2\,, \\ \operatorname{rank} H_1^-(\Sigma) &= \operatorname{rank} H_1(\mathcal{C}) + \#(\mathfrak{b}) - 2 \\ &= -\chi(\mathcal{C}) + \#(\mathfrak{b})\,. \end{aligned} \tag{A.8}$$

However, we can do better than describing ranks. Let

$$(C_\bullet, \partial_\bullet) = \left[0 \to C_2 \xrightarrow{\partial_2} C_1 \xrightarrow{\partial_1} C_0 \to 0\right] \tag{A.9}$$

denote the cellular chain complex corresponding to $D_\bullet$ (so that $C_i = \mathbb{Z}\langle D_i \rangle$), whose homology is $H_\bullet(C_\bullet, \partial_\bullet) = H_\bullet(\Sigma)$. Let

$$(C^-_\bullet, \partial^-_\bullet) = \left[0 \to C^-_2 \xrightarrow{\partial^-_2} C^-_1 \xrightarrow{\partial^-_1} C^-_0 \to 0\right] \tag{A.10}$$

be the chain complex obtained by applying $\ker P_+$ to the groups C_i and specializing the boundary maps, as in Lemma 9. (This makes sense because the deck transformation σ commutes with $\partial_\bullet$.) Notice that C^-_2 and C^-_1 are generated by "odd lifts" of 2-cells and 1-cells

$$C^-_2 = \mathbb{Z}\langle \ell^-(f_i) \rangle\,, \qquad C^-_1 = \mathbb{Z}\langle \ell^-(e_i) \rangle\,, \tag{A.11}$$

where $\ell^-(f_i) := f_i^+ - f_i^-$ and $\ell^-(e_i) := e_i^+ - e_i^-$. It follows from Lemma 9 that $H_1^-(\Sigma) \simeq \ker \partial^-_1 / \operatorname{im} \partial^-_2$ (modulo 2-torsion). More precisely, we have

Lemma 11 *Suppose Σ and $\mathcal{C}$ are both connected. With a cell decomposition and notation as above:*

a) $H_1^-(\Sigma) = \ker \partial^-_1 / (\operatorname{im} \partial_2 \cap C^-_1)$ is generated by 1-cycles $\sum_i a_i \ell^-(e_i)$ formed from odd lifts.

b) If the 0-cells consist entirely of branch points ($D_0 = \mathfrak{b}$) then every odd lift $\ell^-(e_i)$ is automatically closed, so $H_1^-(\Sigma)$ is generated by the $\ell^-(e_i)$.

c) In (b), the intersection product $\langle \ell^-(e_i), \ell^-(e_j) \rangle$ equals the number of common endpoints of e_i and e_j, counted with orientation (which is determined by the choice of lifts made in defining ℓ^-).

Proof. First observe that ∂^-_2 is injective. Indeed, the kernel of ∂_2 is generated by the fundamental class $[\Sigma] = \sum_i (f_i^+ + f_i^-)$, which is even; so if $\partial_2 \beta = 0$, then $\beta = c[\Sigma]$, and $P_-\beta = c\, P_-[\Sigma] = 0$; and if $\beta \in C^-_2$ as well then $P_-\beta = P_+\beta = 0$ imply (Lemma 6(b)) $\beta = 0$. It is also useful to note that $C^\pm_2 = \ker P_\mp|_{C_2} = \operatorname{im} P_\pm|_{C_2}$, and similarly for C_1.

Now, as in the proof of Lemma 9, we generally have that $H_1^-(\Sigma) = \ker P_+ \big| (\ker \partial_1 / \operatorname{im} \partial_2)$ is generated by classes $[\alpha]$ for $\alpha \in C_1$ such that $\partial_1 \alpha = 0$ and $P_+\alpha = \partial_2 \beta$ for some $\beta \in C_2$. Then $0 = P_-P_+\alpha = P_-\partial_2\beta = \partial_2(P_-\beta)$. By injectivity of ∂^-_2 this implies $P_-\beta = 0$, and since $\ker P_-|_{C_2} = \operatorname{im} P_+|_{C_2}$ we get $\beta = P_+\gamma$ for some $\gamma \in C_2$. Setting $\alpha' := \alpha - \partial_2\gamma$, we find $[\alpha] = [\alpha']$ and $P_+\alpha' = 0$. Therefore, $H_1^-(\Sigma)$ is equally well generated by classes $[\alpha]$ for $\alpha \in \ker \partial_1$ such that $P_+\alpha = 0$, and we obtain the identity in part (a).

For part (b), notice that if $D_0 = \mathfrak{b}$ consists entirely of branch points then $\partial_1(e_i^+) = \partial_1(e_i^-)$ for all 1-cells e_i. Therefore $\partial_1\, \ell^-(e_i) = 0$. Part (c) follows by noting that the only intersections of $\ell^-(e_i)$ and $\ell^-(e_j)$ can occur at branch points, and are simple. □

There is one situation not strictly covered by Lemma 11. If $\mathcal{C}$ is connected but Σ is not, then $\Sigma \simeq \mathcal{C}^+ \sqcup \mathcal{C}^-$ is a trivial, disconnected double cover (here the sheets $\mathcal{C}^\pm$ are identical

copies of $\mathcal{C}$). There can be no branch points. The homology $H_1(\Sigma) = H_1(\mathcal{C}^+) \oplus H_1(\mathcal{C}^-)$ is generated by lifts $\gamma^\pm$ of cycles $\gamma \in H_1(\mathcal{C})$ to the two sheets. Let $\ell^-(\gamma) := \gamma^+ - \gamma^-$ denote the odd lift. Than it is easy to see from the decomposition $H_1(\Sigma) = H_1(\mathcal{C}^+) \oplus H_1(\mathcal{C}^-)$ that

Lemma 12 *If $\Sigma \simeq \mathcal{C}^+ \sqcup \mathcal{C}^-$ is a disconnected double cover, then*
a) $\ell^- : H_1(\mathcal{C}) \to H_1^-(\Sigma)$ is an isomorphism; and
b) $\langle \ell^-(\gamma), \ell^-(\gamma') \rangle = 2\langle \gamma, \gamma' \rangle$.

One of the main lessons of Lemmas 11–12 is that for any cover $\Sigma \xrightarrow{\pi} \mathcal{C}$ the homology $H_1^-(\Sigma)$ can be represented by (sums of) curves $\gamma \subset \Sigma$ that are fixed set-wise by σ, such that $\sigma(\gamma) = \overline{\gamma}$ merely flips orientation. These curves are odd lifts $\ell^-(*)$ of either closed curves on $\mathcal{C}$ or edges that connect branch points in a cell decomposition of $\mathcal{C}$. Such generators of $H_1^-(\Sigma)$ are "manifestly" odd. We apply this lesson momentarily.

A.4 Cutting and gluing

A final introductory observation concerns a basic cut-and-glue operation and its odd analogue.

Let S be a closed oriented surface and $\{\mu_i\}_{i=1}^r$ a collection of closed non-intersecting curves on S. Form a surface $S_0 = S \backslash (\sqcup_i \mu_i)$ (possibly disconnected) by cutting along the μ_i. Now S_0 has $2r$ circular boundary components, with images $\mu_i^{(0)}, \mu_i^{(1)}$ of the curves μ_i running along the boundaries. Cap off each boundary component of S_0 with a disc to form another closed, oriented surface S', possibly with multiple components (Figure 29).

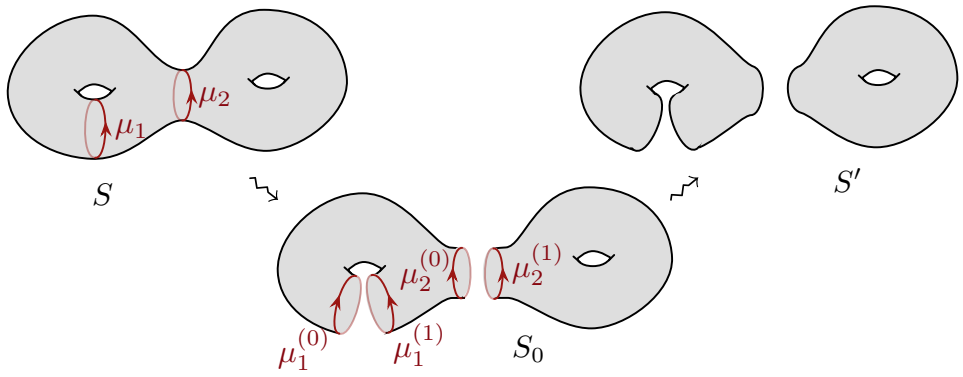

Figure 29. Cutting S along curves μ_i to form S_0, then filling in the boundaries of S_0 with discs to form the closed surface S'. In this case $G = \mathbb{Z}\langle \mu_1 \rangle \subset H_1(S) \simeq \mathbb{Z}^4$ and $H_1(S') = H_1(S)//G \simeq \mathbb{Z}^2$.

Lemma 13 *Let $G \subset H_1(S)$ be the subgroup of $H_1(S)$ generated by the classes of the μ_i's. Let $K = \ker \langle G, * \rangle|_{H_1(S)}$ be the subgroup of elements $H_1(S)$ that have zero intersection with all the μ_i. Then there is a short exact sequence*

$$0 \to G \xrightarrow{i} K \xrightarrow{q} H_1(S') \to 0 \tag{A.12}$$

that identifies $H_1(S') \simeq K/G =: H_1(S)//G$ as a lattice symplectic quotient. In particular, the intersection form on $H_1(S')$ is induced by the quotient from the form on $H_1(S)$.

Proof. First observe that since the μ_i's are non-intersecting, the intersection form vanishes on G, $\langle G, G\rangle = 0$. This it makes sense to define $K = \ker \langle G, *\rangle|_{H_1(S)}$, and $i : G \hookrightarrow K$ is just the inclusion. Exactness of the rest of (A.12) can be derived by comparing long exact sequences in relative homology for the pairs (S, S_0) and (S', S_0). However, it is useful to take a more concrete approach.

The map q is defined as follows. Any homology class $[\gamma] \in K \subset H_1(S)$ can be represented by a curve (or sum of curves) γ that does not intersect the μ_i's. Thus γ can be naturally included in S_0 and (viewing S_0 as a subspace of S') as a curve $\gamma' \subset S'$. We set $q([\gamma]) := [\gamma']$. The only potential ambiguity in this definition comes from the μ_i themselves, which may be included either as $\mu_i^{(0)}$ or $\mu_i^{(1)}$ in S_0 and then in S'. But in S' the curves $\mu_i^{(0)}, \mu_i^{(1)}$ bound discs, so they are contractible, and $q([\mu_i]) = 0$ unambiguously. Thus q is well defined and its kernel includes all of G. Moreover, it preserves the intersection form.

Conversely, any $[\gamma'] \in H_1(S')$ can be represented by a curve $\gamma' \subset S_0 \subset S'$ (*i.e.* a curve in the complement of the discs that filled in the holes of S_0). Viewing S_0 as a subspace of S, we can thus include γ' as a curve $\gamma \subset S$, and try to define $q^{-1}([\gamma']) = [\gamma] \in K \subset H_1(S)$. There is ambiguity in this definition coming from how the representative γ' is chosen to wind around the discs in $S'\backslash S_0$; this ambiguity is precisely G. Therefore, we get a map $q^{-1} : H_1(S') \to K/G$ that provides an inverse to q, and demonstrates that $q : K/G \xrightarrow{\sim} H_1(S')$ is an isomorphism. □

We actually want a generalization of (A.12) to odd homology. To this end, suppose that we perform a basic cut-and-glue operation $\mathcal{C} \rightsquigarrow \mathcal{C}_0 \rightsquigarrow \mathcal{C}'$ along some non-intersecting closed curves $\overline{\mu}_i \subset \mathcal{C}$. Suppose that the $\overline{\mu}_i$ are chosen to lie in the complement of the branching locus $\mathfrak{b}$ of a double cover $\Sigma \xrightarrow{\pi} \mathcal{C}$. Then the cut-and-glue operation lifts to Σ. We first take all lifts of the $\overline{\mu}_i$ to Σ, noting that a given $\overline{\mu}_i$ may have a single lift μ_i if $\pi^{-1}(\mu_i)$ is connected and two lifts $\mu_i^{\pm}$ otherwise. We then cut Σ along the $\mu_i^{\pm}, \mu_i$ to form Σ_0, and we fill in the boundaries of Σ_0 with discs to form Σ'. There are induced covering maps $\Sigma_0 \xrightarrow{\pi} \mathcal{C}_0$ and $\Sigma' \xrightarrow{\pi} \mathcal{C}'$, such that each disc in $\mathcal{C}'\backslash\mathcal{C}_0$ contains a new branch point and is covered by a single disc in $\Sigma'\backslash\Sigma_0$ if and only if the corresponding cutting curve $\overline{\mu}_i$ has a single connected lift μ_i (otherwise a disc in $\mathcal{C}'\backslash\mathcal{C}_0$ is covered by two disconnected discs in $\Sigma'\backslash\Sigma_0$). We now have

Lemma 14 *Let $G \subset H_1(\Sigma)$ denote the subgroup generated by the lifts $\mu_i^{\pm}, \mu_i$, and $K := \ker \langle G, *\rangle|_{H_1(\Sigma)}$. Then G and K are fixed by the deck transformation σ_*, and we can take odd parts $G^- = \ker P_+|_G = G \cap H_1^-(\Sigma)$ and $K^- = \ker P_+|_K = K \cap H_1^-(\Sigma) = \ker\langle G^-, *\rangle|_{H_1^-(\Sigma)}$. There is a complex*

$$0 \to G^- \xrightarrow{i^-} K^- \xrightarrow{q^-} H_1^-(\Sigma') \to 0 \tag{A.13}$$

with vanishing homology at all but the last spot, where the homology is 2-torsion. The maps i^-, q^- preserve the intersection form. Therefore there is an injection $H_1^-(\Sigma)//G^- = K^-/G^- \xhookrightarrow{q^-} H_1^-(\Sigma')$ with finite (2-torsion) cokernel, which preserves the intersection form.

If every cutting curve $\overline{\mu}_i \subset \mathcal{C}$ has two distinct lifts $\mu_i^{\pm}$ to Σ, then (A.13) is exact, and $q^- : K^-/G^- \hookrightarrow H_1^-(\Sigma')$ is an isomorphism.

Proof. First observe that σ_* preserves G because $\sigma_*\mu_i^{\pm} = \mu_i^{\mp}$ and $\sigma_*\mu_i = \mu_i$; and σ_* preserves K because it preserves the intersection form. So $G^- := \ker P_+|_G$ and $K^- := \ker P_+|_K$ make sense. Also, K^- coincides with $\ker \langle G^-, * \rangle|_{H_1^-(\Sigma)}$: inclusion $\ker P_+|_K \subset \ker \langle G^-, * \rangle|_{H_1^-(\Sigma)}$ is obvious; conversely if $\alpha \in \ker \langle G^-, * \rangle|_{H_1^-(\Sigma)}$ then $\langle \mu_i^+ - \mu_i^-, \alpha \rangle = 0 \Rightarrow \langle \mu_i^+, \alpha \rangle = \langle \mu_i^-, \alpha \rangle = \langle \sigma\mu_i^-, \sigma\alpha \rangle = -\langle \mu_i^+, \alpha \rangle \Rightarrow \langle \mu_i^{\pm}, \alpha \rangle = 0$, and $\langle \mu_i, \alpha \rangle = 0$ because μ_i is even and α is odd; so $\alpha \in \ker \langle G, * \rangle|_{H_1(\Sigma)} = K$, hence $\ker \langle G^-, * \rangle|_{H_1^-(\Sigma)} \subset \ker P_+|_K$.

Consider the exact sequence $0 \to G \xrightarrow{i} K \xrightarrow{q} H_1(\Sigma') \to 0$ from Lemma 13. Since σ_* commutes with i and q, we can take $\ker P_+$ of all groups and apply Lemma 9(c) to obtain the complex (A.13), with vanishing homology except perhaps at the last spot, where homology must be 2-torsion. The restricted maps i^- and q^- must preserve the intersection form, because i and q do.

In the case that every $\overline{\mu}_i \subset \mathcal{C}$ has two distinct lifts, the discs in $\Sigma' \backslash \Sigma_0$ contain no branch points. Thus (as at the end of Appendix A.3) every element of $H_1^-(\Sigma')$ can be represented by a sum of curves $\gamma' \subset \Sigma_0 \subset \Sigma'$ that is manifestly odd, *i.e.* σ fixed γ' setwise while reversing orientation. Then the inverse map q^{-1} of Lemma 13 sends $[\gamma']$ to an element $[\gamma] \in K/G$ that is *also* manifestly odd, and its restriction $q^{-1}|_{H_1^-(\Sigma')} : H_1^-(\Sigma') \to K^-/G^-$ provides an inverse to q^- in (A.13), guaranteeing that $q^- : K^-/G^- \hookrightarrow H_1^-(\Sigma')$ is actually an isomorphism.

(If there are branch points on the discs in $\Sigma' \backslash \Sigma_0$ this can fail. In order to represent all $[\gamma'] \in H_1^-(\Sigma')$ by curves γ' that are manifestly odd, the curves may have to go through the branch points — and cannot be deformed into $\Sigma_0 \subset \Sigma$ while keeping them manifestly odd. Thus while the inverse $q^{-1}|_{H_1^-(\Sigma')} : H_1^-(\Sigma') \to K/G$ still exists, its image may not be completely odd. This is the source of the 2-torsion in (A.13).) $\square$

B Reconstructing framed flat connections

Let M by a framed, triangulated 3-manifold with $\pi_1(\mathcal{C}_{\text{small}})$ abelian. We briefly explain how to uniquely reconstruct a framed flat connection $\mathcal{A} \in \mathcal{P}[\mathcal{C}] \subset \mathcal{X}[\mathcal{C}]$ given functions $x_{\mathfrak{p}}$ satisfying the conditions on the RHS of (4.14), proving part (d) of Proposition 2 in the process.

The basic idea (following [44] and precursors, *e.g.* [49]) is to use the functions to construct a distinguished set of projective bases (*i.e.* frames) for the fiber of the bundle $E \to \mathcal{C}$ at various points of $\mathcal{C}$, together with maps between these bases corresponding to the parallel transport of $\mathcal{A}$. There are several steps.

Start with a connected component of the big boundary $\mathcal{C}_{\text{big}}$, and a hexagon f in the triangulation $\mathbf{t}_{2\text{d}}$. We assume that the flat connection is trivialized in the interior f° of f. Choose three arbitrary distinct lines a, b, c in the fiber of E over f°, and make them framing lines at the three small edges of f (this choice fixes gauge redundancy). Construct projective bases $b(p)$ for the fiber of E over six points $p \in f^\circ$ as in Figure 30, as follows. For the element in $b(p)$ take any vector v_1 in the framing line on the small edge of f closest to p. For the second element take a vector v_2 in the framing line on the small edge of f second-closed to p, normalized that $v_1 \pm v_2$ lies in the third framing line. The sign is specified by orientation, as

in Figure 30. The $PGL(2)$ transformations among the six bases in f are given by matrices $S = \left(\begin{smallmatrix} 0 & -1 \\ 1 & 0 \end{smallmatrix}\right)$ and $T = \left(\begin{smallmatrix} 1 & 0 \\ 1 & 1 \end{smallmatrix}\right)$, as in the Figure, obeying $(ST)^3 = 1$ in $PGL(2)$.

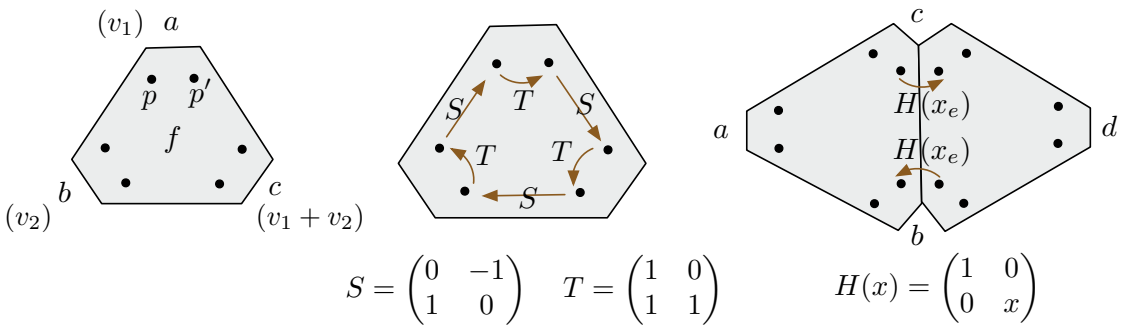

Figure 30. Left: projective bases assigned to six points in a hexagon; for example at p the basis is (v_1, v_2) with $v_1 \in a$, $v_2 \in b$, $v_1 + v_2 \in c$; while at p' the basis is (v_1, w_2) with $v_1 \in a$, $w_2 \in b$, $v_1 - w_2 \in c$. Middle: $PGL(2)$ transformations among bases in any hexagon. Right: $PGL(2)$ transformations across big edges.

Proceed to assign three framing lines and six projective bases to the rest of the hexagons in the connected component of $\mathcal{C}_{\text{big}}$, subject to the following rules:
1) framing lines at adjacent small edges (bounding a hole in $\mathcal{C}_{\text{big}}$) must agree;
2) the connection is trivialized in the interior of every hexagon, and the six bases there are constructed the same way they were for f, with the same $PGL(2)$ relations;
3) parallel transport across a big edge e is given by $H(x_e) = \left(\begin{smallmatrix} 1 & 0 \\ 0 & x_e \end{smallmatrix}\right) \in PGL(2)$, where x_e is the edge function .
Following these rules, there are no more arbitrary choices to be made, and we reconstruct a framed flat connection on $\mathcal{C}_{\text{big}}$ — with framing lines at the holes of $\mathcal{C}_{\text{big}}$. We repeat for every connected component of $\mathcal{C}_{\text{big}}$.

Next, we extend the framed flat connection over the small boundary. Notice that the holonomy on a clockwise path around any hole on the big boundary takes the form

$$\text{Hol}(\textit{hole}) = \begin{pmatrix} 1 & 0 \\ * & \prod_{e \text{ at hole}} x_e^{-1} \end{pmatrix} . \tag{B.1}$$

This preserves the framing line on the boundary of the hole, rescaling it with (squared) eigenvalue $\prod_e x_e$. If the hole is to be filled in with a small disc, then $\prod_e x_e = 1$ by (4.14), so we can uniquely extend the framed flat connection over a *punctured* disc with unipotent holonomy at the puncture (as required for $\mathcal{X}[\mathcal{C}]$).

If instead a pair of holes h, h' are connected by a small annulus, we choose two hexagons f, f' adjacent to the holes and a path $\mathfrak{p}_\tau$ running from f to f'. Let $p \in f$, $p' \in f'$ be points with projective bases on the two sides of the annulus, as in Figure 31. Let M_λ, M'_λ be the $PGL(2)$ holonomies around the ends of the annulus, with basepoints at p, p', running clockwise (resp., counter-clockwise) from the viewpoint of h (resp., h'). The holonomies are

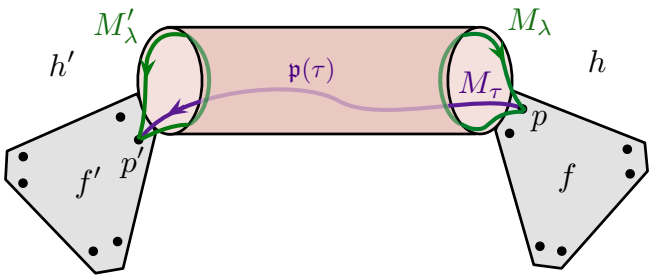

Figure 31. Reconstructing a framed flat connection on an annulus: the parallel transports must satisfy $M_\lambda = M_\tau^{-1} M'_\lambda M_\tau$.

fixed to be

$$M_\lambda = \begin{pmatrix} 1 & 0 \\ a & x_\lambda^{-1} \end{pmatrix}, \quad M'_\lambda = \begin{pmatrix} 1 & 0 \\ a' & x_\lambda^{-1} \end{pmatrix}, \qquad x_\lambda = \prod_{e \text{ at } h} x_e = \prod_{e' \text{ at } h'} x_{e'}^{-1} \tag{B.2}$$

for some (determined) a, a'. The function $x_\tau = x_{\mathfrak{p}(\tau)}$ fixes the eigenvalue of the $PGL(2)$ transformation M_τ between the bases at p and p', given by parallel transport along $\mathfrak{p}_\tau$,

$$M_\tau = \begin{pmatrix} 1 & 0 \\ t & x_\tau^{-1} \end{pmatrix} \tag{B.3}$$

for some undetermined t. In order to identify the framing lines at the two ends of the annulus and extend the framed flat connection over the annulus, it suffices to require $M_\lambda = M_\tau^{-1} M'_\lambda M_\tau$, which uniquely determines $t = (a' - a x_\tau^{-1})/(1 - x_\lambda^{-1})$ as long as $x_\lambda \neq 1$.[22]

This completes the reconstruction of $\mathcal{A}$ over any connected component of $\mathcal{C}$ containing both big and small parts. In particular, the $PGL(2)$ holonomy along any closed path may be obtained by combining the "traffic rules" of Figure 30 for the big boundary with the transport M_τ along annuli. Triviality of the holonomy along any contractible path is ensured by the local identities $(ST)^3 = H(x)SH(x)S = M_\lambda^{-1} M_\tau^{-1} M'_\lambda M_\tau = 1$.

Finally, on an isolated small sphere the framed flat connection is trivial. (The connection itself is trivial, and its residual gauge symmetry can be used to trivialize the framing data.) For a small torus, if we are given a choice of A and B cycles with (squared) holonomy eigenvalues $(x_\alpha, x_\beta) \neq (1, 1)$, we may uniquely reconstruct the commuting holonomy matrices (modulo $PGL(2)$ gauge equivalence) as $M_\alpha = H(x_\alpha^{-1})$, $M_\beta = H(x_\beta^{-1})$. We take the framing to be the unique eigenline with (squared) eigenvalues (x_α, x_β).

[22] Alternatively, one can think of M_τ as the $PGL(2)$ gauge transformation required to identify the holonomies M_λ and M'_λ at the two ends of an annulus. Determining M_τ fixes a relative gauge ambiguity between components of $\mathcal{C}_{\text{big}}$ that are connected by annuli.

References

[1] E. Witten, *Monopoles and Four-Manifolds*, *Math. Res. Let.* **1** (1994) 769–796, [hep-th/9411102v1].

[2] D. Gaiotto, *N=2 dualities*, *JHEP* **1208** (2012) 034, [arXiv:0904.2715].

[3] D. Gaiotto, G. W. Moore, and A. Neitzke, *Four-dimensional wall-crossing via three-dimensional field theory*, *Comm. Math. Phys.* **299** (2010), no. 1 163–224, [arXiv:0807.4723].

[4] D. Gaiotto, G. W. Moore, and A. Neitzke, *Wall-crossing, Hitchin Systems, and the WKB Approximation*, *Adv. Math.* **234** (2013) 239–403, [arXiv:0907.3987].

[5] D. Gaiotto, G. W. Moore, and A. Neitzke, *Wall-Crossing in Coupled 2d-4d Systems*, *JHEP* **1212** (2012) 082, [arXiv:1103.2598].

[6] D. Gaiotto, G. W. Moore, and A. Neitzke, *Spectral networks*, *Ann. Henri Poincare* **14** (2013), no. 7 1643–1731, [arXiv:1204.4824].

[7] T. Dimofte, S. Gukov, and L. Hollands, *Vortex Counting and Lagrangian 3-manifolds*, arXiv:1006.0977.

[8] Y. Terashima and M. Yamazaki, *SL(2,R) Chern-Simons, Liouville, and Gauge Theory on Duality Walls*, *JHEP* **1108** (2011) 135, [arXiv:1103.5748].

[9] T. Dimofte and S. Gukov, *Chern-Simons Theory and S-duality*, arXiv:1106.4550.

[10] T. Dimofte, D. Gaiotto, and S. Gukov, *Gauge Theories Labelled by Three-Manifolds*, arXiv:1108.4389.

[11] W. Thurston, *The Geometry and Topology of Three-Manifolds*, *Lecture notes at Princeton University* (1980).

[12] W. D. Neumann and D. Zagier, *Volumes of hyperbolic three-manifolds*, *Topology* **24** (1985), no. 3 307–332.

[13] D. W. Boyd, F. Rodriguez-Villegas, and N. M. Dunfield, *Mahler's Measure and the Dilogarithm (II)*, *Canad. J. Math.* **54** (2002), no. 3 468–492, [math/0308041v2].

[14] A. Champanerkar, *A-Polynomial and Bloch Invariants of Hyperbolic 3-Manifolds*, *Ph.D. Thesis, Columbia University* (2003).

[15] K. Hikami, *Generalized Volume Conjecture and the A-Polynomials - the Neumann-Zagier Potential Function as a Classical Limit of Quantum Invariant*, *J. Geom. Phys.* **57** (2007), no. 9 1895–1940, [math/0604094v1].

[16] T. Dimofte, S. Gukov, J. Lenells, and D. Zagier, *Exact Results for Perturbative Chern-Simons Theory with Complex Gauge Group*, *Comm. Num. Thy. and Phys.* **3** (2009), no. 2 363–443, [arXiv:0903.2472].

[17] T. Dimofte, *Quantum Riemann Surfaces in Chern-Simons Theory*, arXiv:1102.4847.

[18] J. E. Andersen and R. Kashaev, *A TQFT from quantum Teichmüller theory*, arXiv:1109.6295.

[19] T. D. Dimofte and S. Garoufalidis, *The quantum content of the gluing equations*, *arXiv* **math.GT** (2012) [arXiv:1202.6268].

[20] S. Garoufalidis, *The 3D index of an ideal triangulation and angle structures*, arXiv:1208.1663.

[21] S. Gukov, *Three-Dimensional Quantum Gravity, Chern-Simons Theory, and the A-Polynomial*, *Commun. Math. Phys.* **255** (2005), no. 3 577–627, [hep-th/0306165v1].

[22] S. Garoufalidis, *On the Characteristic and Deformation Varieties of a Knot*, *Geom. Topol. Monogr.* **7** (2004) 291–304, [math/0306230v4].

[23] C. Frohman, R. Gelca, and W. Lofaro, *The A-Polynomial From the Noncommutative Viewpoint*, *Trans. Amer. Math. Soc.* **354** (2002), no. 2 735–747, [math/9812048v1].

[24] R. Dijkgraaf and H. Fuji, *The Volume Conjecture and Topological Strings*, *Fortschr. Phys.* **57** (2009), no. 9 825–856, [arXiv:0903.2084].

[25] S. Gukov and P. Sułkowski, *A-polynomial, B-model, and Quantization*, arXiv:1108.0002.

[26] G. Borot and B. Eynard, *All-order asymptotics of hyperbolic knot invariants from non-perturbative topological recursion of A-polynomials*, arXiv:1205.2261.

[27] T. Dimofte, D. Gaiotto, and S. Gukov, *3-Manifolds and 3d Indices*, arXiv:1112.5179.

[28] E. Witten, *Fivebranes and Knots*, *Quantum Topol.* **3** (2012), no. 1 1–137, [arXiv:1101.3216].

[29] W. Neumann, *Combinatorics of Triangulations and the Chern-Simons Invariant for Hyperbolic 3-Manifolds*, *in Topology '90, Ohio State Univ. Math. Res. Inst. Publ.* **1** (1992).

[30] C. K. Zickert, *The Volume and Chern-Simons Invariant of a Representation*, *Duke Math. J.* **150** (2009), no. 3 489–532, [arXiv:0710.2049].

[31] T. Dimofte, M. Gabella, and A. B. Goncharov, *K-Decompositions and 3d Gauge Theories*, arXiv:1301.0192.

[32] T. Dimofte, D. Gaiotto, and R. van der Veen, *RG Domain Walls and Hybrid Triangulations*, arXiv:1304.6721.

[33] A. Guilloux, *Representations of 3-manifold groups in PGL(n,C) and their restriction to the boundary*, arXiv:1310.2907.

[34] F. Bonahon, *A Schlafli-type formula for convex cores of hyperbolic 3-manifolds*, *J. Diff. Geom.* **50** (1998), no. 1 25–58, [dg-ga/9704017v1].

[35] N. Bergeron, E. Falbel, and A. Guilloux, *Tetrahedra of flags, volume and homology of SL(3)*, arXiv:1101.2742.

[36] S. Garoufalidis and C. K. Zickert, *The symplectic properties of the PGL(n,C)-gluing equations*, arXiv:1310.2497.

[37] D. Gaiotto, G. W. Moore, and A. Neitzke, *Spectral Networks and Snakes*, arXiv:1209.0866.

[38] A. B. Goncharov and M. Kontsevich *Unpublished work.*

[39] M. F. Atiyah and R. Bott, *The Yang-Mills equations over Riemann surfaces*, *Philos. Trans. Roy. Soc. London Ser. A* **308** (1983), no. 1505 523–615.

[40] L. Hollands and A. Neitzke, *Spectral networks and Fenchel-Nielsen coordinates*, arXiv:1312.2979.

[41] N. Nekrasov, A. Rosly, and S. Shatashvili, *Darboux coordinates, Yang-Yang functional, and gauge theory*, *Nucl. Phys. Proc. Suppl.* **216** (2011) 69–93, [arXiv:1103.3919].

[42] Y. Kabaya, *Parametrization of PSL(2,C)-representations of surface groups*, arXiv:1110.6674.

[43] F. Luo, S. Tillmann, and T. Yang, *Thurston's Spinning Construction and Solutions to the Hyperbolic Gluing Equations for Closed Hyperbolic 3-Manifolds*, arXiv:1004.2992.

[44] V. V. Fock and A. B. Goncharov, *Moduli spaces of local systems and higher Teichmuller theory*, *Publ. Math. Inst. Hautes Etudes Sci.* **103** (2006) 1–211, [math/0311149v4].

[45] A. Hatcher, *Algebraic Topology*, *Cambridge Univ. Press* (2002).

[46] S. Matveev, *Algorithmic topology and classification of 3-manifolds*, *Algorithms and Computation in Mathematics, Springer* **9** (2007) xiv+492.

[47] S. Tillmann, *Normal surfaces in topologically finite 3-manifolds*, *Enseign. Math. (2)* **54** (2008), no. 3-4 329–380, [math/0406271v3].

[48] D. Futer and F. Guéritaud, *From angled triangulations to hyperbolic structures*, *Contemp. Math.* **541** (2011) 159–182, [arXiv:1004.0440].

[49] V. V. Fock, *Dual Teichmüller spaces*, dg-ga/9702018v3.